

Coalgebra and Modal Logic: a first introduction

Yde Venema*

©YV 2026

DO NOT DISTRIBUTE

Abstract

These notes give a brief introduction to the theory of universal coalgebra and coalgebraic modal logic.

*Institute for Logic, Language and Computation, University of Amsterdam, Science Park 107, NL-1098XG Amsterdam. E-mail: y.venema@uva.nl.

Contents

1	Introduction	1-1
2	Final Coalgebras and Coinduction	2-1
3	Bisimilarity and Behavioural Equivalence	3-1
4	Covarieties	4-1
5	Coalgebraic modalities via relation lifting	5-1
6	Coalgebraic modalities via predicate liftings	6-1
7	One-step logic	7-1
8	Soundness and completeness	8-1
9	Coalgebraic fixpoint logic	9-1
A	Appendix: Basic mathematical definitions	A-1
B	Appendix: The Category Set and its Functors	B-1

1 Introduction

Starting from concrete examples, this chapter introduces (set-based) coalgebras, together with some of the most important coalgebraic concepts, including coinduction, behavioural equivalence and bisimilarity. We then give a first discussion of the relation between coalgebra and modal logic, and we give some examples of coalgebraic modal logics. In the final section we provide some pointers to the literature on coalgebra and modal logic.

1.1 State-based evolving systems

Example 1.1 Perhaps the simplest example of a computational process is the following black-box machine with two buttons, h and n . If we press the h -button, the machine displays some value from a data set C . No matter how many times we press the h -button, this value remains the same. Each time we push the n -button, however, we may observe a different value by pushing the h -button.

A natural way to formally describe this device is as a set S of internal states (that are not visible to the user), together with two maps

$$\begin{aligned} h &: S \rightarrow C \\ n &: S \rightarrow S, \end{aligned}$$

where $h(s)$ indicates the observable value at state s , and $n : S \rightarrow S$ is a function mapping a state $s \in S$ to its unique next state.

All that we may observe of a state s , is the *stream*

$$h(s) \cdot h(n(s)) \cdot h(n(n(s))) \cdot h(n^3(s)) \cdots$$

of data. This stream will be called the *behaviour* of s in $\mathbb{S} = (S, h, n)$, notation: $\text{beh}_{\mathbb{S}}(s)$.

Two states s and s' in two black boxes can then be called *behaviourally equivalent*, notation: $\mathbb{S}, s \simeq \mathbb{S}', s'$, if they display the same behaviour, that is, if $\text{beh}_{\mathbb{S}}(s) = \text{beh}_{\mathbb{S}'}(s')$.

Example 1.2 A *deterministic finite state automaton* or *DFA* over an alphabet C is a triple $\mathbb{A} = (A, \delta, F)$, where A is a finite set of states, $\delta : A \times C \rightarrow A$ is the *transition map* of $\mathbb{A}$, and $F \subseteq A$ is the set of *accepting states* of the automaton. In contrast to the usual presentation we do not take the initial state of a DFA to be part of its structure; rather, we will consider *initialised versions* of DFAs, where we explicitly single out an initial state for the automaton.

Let C^* denote the set of finite words over C , then we may extend δ to a map $\widehat{\delta} : A \times C^* \rightarrow A$ as follows:

$$\begin{aligned} \widehat{\delta}(a, \varepsilon) &:= a \\ \widehat{\delta}(a, cw) &:= \widehat{\delta}(\delta(a, c), w). \end{aligned}$$

We define, for a state $a \in A$,

$$L_{\mathbb{A}}(a) := \{w \in C^* \mid \widehat{\delta}(a, w) \in F\},$$

as the *language* accepted by $\mathbb{A}$, initialized at a . Two initialized automata $(\mathbb{A}, a)$ and $(\mathbb{A}', a')$ are (*language*) *equivalent* if they accept exactly the same words, that is, if $L_{\mathbb{A}}(a) = L_{\mathbb{A}'}(a')$.

Example 1.3 The key structures featuring in the semantics of *modal logic* are Kripke frames and Kripke models. A *Kripke frame* is a pair (S, R) consisting of a set S of objects called *states*, *points* or *worlds*, and an *accessibility relation* $R \subseteq S \times S$. A *Kripke model* is a triple $\mathbb{S} = (S, R, V)$ such that (S, R) is a Kripke frame (the *underlying Kripke frame* of the model), and V is a *valuation*, i.e., a map $Q \rightarrow PS$, where Q is some fixed set of *proposition letters*.

A *bisimulation* between two Kripke models $\mathbb{S}$ and $\mathbb{S}'$ is a binary relation $Z \subseteq S \times S'$ such that, for all $s \in S$ and $s' \in S'$ with Zss' the following conditions hold:

- (atom) s and s' satisfy the same proposition letters;
- (forth) for all $t \in S$ such that Rst there is a $t' \in S'$ with $R's't'$ and Ztt' ;
- (back) for all $t' \in S'$ such that $R's't'$ there is a $t \in S$ with Rst and Ztt' .

If there is a bisimulation Z linking s and s' we say that s and s' are *bisimilar*, notation $\mathbb{S}, s \Leftrightarrow \mathbb{S}', s'$ (or $Z : \mathbb{S}, s \Leftrightarrow \mathbb{S}', s'$ if we want to make the bisimulation explicit).

Given a modal language L , we let $Th_{\mathbb{S}}^L(s)$ denote the set of L -formulas that are true at s in $\mathbb{S}$; then two states are called *(L-)equivalent* (notation: $\equiv_L$) if they satisfy the same L -formulas. A theme in the theory of modal logic is to study the relation between equivalence and bisimilarity. A class of models C is said to have the *Hennessy-Milner* property with respect to a language L if $\equiv_L = \Leftrightarrow$ on C .

Example 1.4 The theory of *non-well-founded sets* provides an alternative to the standard axiomatic set theories by allowing sets to contain themselves, or otherwise violate the rule of well-foundedness. More in detail, in non-well-founded set theories, the Foundation Axiom *FA* is replaced by axioms implying its negation. For instance, working with the *anti-foundation axiom* *AFA* we may associate, with each so-called *apg* or *accessible pointed graph* (that is, a directed graph such that every node can be reached via a finite path from a specified *root* of the graph) a *hyperset*, that is, a set that is not necessarily well-founded. And, two apgs yield the *same* set iff they are bisimilar.

Example 1.5 As a final example of a state-based evolving system we mention Markov chains: transition systems that evolve probabilistically. Recall that a *(discrete) probability distribution* on a set S is a map $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$. Formally, a Markov chain can be modelled as a pair (S, σ) , where σ assigns a probability distribution σ_s to each state s .

For a concrete example, think of a gambler wagering €1 on a series of fair coin tosses – this series may be indefinite, or end if the gambler loses his money. This experiment can be modelled by the Markov chain (S, σ) where $S = \{s_n \mid n \in \omega\}$, with state s_n representing the state where the gambler owns € n . For $n > 0$ we have that σ_{s_n} assigns a 0.5 probability to both s_{n-1} and s_{n+1} (and a probability 0 to all other states), while σ_{s_0} assigns a 1.0 probability to s_0 (and a probability 0 to all other states).

1.2 Coalgebras and their morphisms

As we will see now, the structures described in the previous section all are specimens of *coalgebras*. Universal Coalgebra is a theory of state-based evolving systems, formulated in the language of category theory.¹

¹See the appendix for some background definitions on category theory.

Definition 1.6 Given an endofunctor $T : \mathbf{C} \rightarrow \mathbf{C}$ on some category $\mathbf{C}$, a T -coalgebra is a pair (X, ξ) where X is an object in $\mathbf{C}$ and $\xi : X \rightarrow TX$ is an arrow in $\mathbf{C}$. We will sometimes refer to T as the *type* of (X, ξ) . If, for an arrow $f : X' \rightarrow X$, the following diagram commutes:

$$\begin{array}{ccc} X' & \xrightarrow{f} & X \\ \xi' \downarrow & & \downarrow \xi \\ TX' & \xrightarrow{Tf} & TX \end{array} \quad (1)$$

we call f a (*coalgebra*) *morphism* from $\mathbb{X}' = (X', \xi')$ to $\mathbb{X} = (X, \xi)$, and write $f : \mathbb{X}' \rightarrow \mathbb{X}$.

We let $\mathbf{Coalg}_{\mathbf{C}}(T)$ denote the category with T -coalgebras as objects and T -coalgebra morphisms as arrows; the category $\mathbf{C}$ will be called the *base* category of $\mathbf{Coalg}_{\mathbf{C}}(T)$. $\triangleleft$

We will usually (but not always) confine our attention to *systems*, that is, coalgebras over the category $\mathbf{Set}$. Intuitively, a set functor T specifies the one-step dynamics that a system can engage in.

Definition 1.7 A *set functor* is an endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ on the category $\mathbf{Set}$ of sets and functions; given such a set functor T , we will sometimes refer to T -coalgebras as T -*systems*. Given the prominence of these coalgebras we will often simply write $\mathbf{Coalg}(T)$ rather than $\mathbf{Coalg}_{\mathbf{Set}}(T)$ to denote the class of all T -systems. Where $\mathbb{S} = (S, \sigma)$ is a T -system, we refer to S and σ as, respectively, the *carrier* or *state space* and the *transition map* or *coalgebra map* of $\mathbb{S}$. A *pointed* or *initialized* T -system is a triple (S, σ, s) such that (S, σ) is a T -system and $s \in S$. $\triangleleft$

The coalgebraic viewpoint on systems combines wide applicability and mathematical simplicity: since every set functor determines its own type of coalgebra, notions, properties and results of state-based systems can be uniformly explained just in terms of properties of their type functors. This applies to systems as diverse as streams, probabilistic transition systems, automata, Kripke structures and neighbourhood frames. In the appendix we give a list of set functors; here we give a few examples of the associated coalgebras.

Example 1.8 (a) The black boxes of Example 1.1 are systems of the functor type $K_C \times Id$, where K_C is the constant functor associated with the set C , and Id is the identity functor on $\mathbf{Set}$.

(b) Deterministic finite automata (Example 1.2) are systems of type $2 \times Id^C$, where 2 is the set $\{0, 1\}$. To see this, consider a coalgebra $\mathbb{X} = (X, \xi)$ of this type; then ξ determines, for each state $x \in X$, two things: an element $\xi_0(x)$ of the set 2 , specifying whether x is accepting ($\xi_0(x) = 1$) or not ($\xi_0(x) = 0$), and an element $\xi_1(x) \in X^C$, that is, a map $\xi_1(x) : C \rightarrow X$ providing a successor of x in X for each letter $c \in C$. A $2 \times Id^C$ -coalgebra (X, ξ) thus corresponds to the DFA $(X, \xi_0^{-1}(1), \xi_1)$.

(c) Kripke frames are coalgebras for the powerset functor P , whereas Kripke models are coalgebras for the functor $K_{PQ} \times P$. We leave it for the reader to verify that, given two Kripke frames $\mathbb{S}$ and $\mathbb{S}'$, a map $f : S \rightarrow S'$ is a coalgebra morphism iff it is a bounded morphism, that is, for every $s \in S$ it satisfies:

- if $(s, t) \in R$ then $(fs, ft) \in R'$;
- if $(fs, t') \in R'$ then $(s, t) \in R$ for some t such that $ft = t'$.
- (d) Non-well-founded sets can be represented by certain pointed P -systems.
- (e) Markov chains are D -systems, where D is the distribution functor (which assigns, to a set S , a discrete probability distribution on S).
- (f) For every set functor T we will allow the empty T -coalgebra $(\emptyset, \emptyset)$.

1.3 Final coalgebras and coinduction

For many coalgebra types T one may associate with an arbitrary state s in an arbitrary T -coalgebra $\mathbb{S}$, a natural notion of *behaviour*. This can often be formalised by defining a *behaviour map* and proving that this map is the unique coalgebra morphism from $\mathbb{S}$ to some *final* or *terminal* coalgebra $\mathbb{Z}$ of type T .

Definition 1.9 Let T be an endofunctor on some category $\mathbf{C}$. A T -coalgebra $\mathbb{Z} = (Z, \zeta)$ is *final* or *terminal* if it is a final object in the category $\mathbf{Coalg}_{\mathbf{C}}(T)$; that is, if for every T -coalgebra $\mathbb{X} = (X, \xi)$ there is a *unique* morphism from $\mathbb{X}$ to $\mathbb{Z}$; this morphism will be denoted as $\text{beh}_{\mathbb{X}} : \mathbb{X} \rightarrow \mathbb{Z}$. $\triangleleft$

Note that final coalgebras, when they exist, are unique modulo isomorphism. For this reason we will often speak of *the* final T -coalgebra of a functor T .

Example 1.10 A *stream* over a set C is a map $\alpha : \omega \rightarrow C$ (where ω is the set of natural numbers). We may turn the set C^ω of C -streams into a $K_C \times Id$ -coalgebra itself by endowing it with a coalgebra map $\gamma := (\mathbf{h}, \mathbf{t}) : C^\omega \rightarrow C \times C^\omega$. Here we define the maps $\mathbf{h} : C^\omega \rightarrow C$ and $\mathbf{t} : C^\omega \rightarrow C^\omega$ by putting, for an arbitrary C -stream $\alpha : \omega \rightarrow C$:

$$\begin{aligned} \mathbf{h}(\alpha) &:= \alpha(0) \\ \mathbf{t}(\alpha) &:= \lambda n. \alpha(n+1). \end{aligned}$$

That is, the coalgebra map γ splits an infinite C -stream $c_0c_1c_2\cdots$ into its *head* c_0 and its *tail* $c_1c_2c_3\cdots$.

It is then not very hard to prove that the *stream coalgebra* (C^ω, γ) is a final coalgebra for the functor $K_C \times Id$: this boils down to showing that, for an arbitrary ‘black box machine’ $\mathbb{S} = (S, h, n)$, the behaviour map $\text{beh} : S \rightarrow C^\omega$ is the *unique* coalgebra morphism from $\mathbb{S}$ to (C^ω, γ) .

Example 1.11 For a second example, fix an alphabet C and define a (C) -*language* to be any set of finite words. Further on we will see that we can endow the collection $P(C^*)$ of C -languages with a very natural coalgebra map for the functor $2 \times Id^C$ of deterministic finite C -automata, and prove that the resulting structure is in fact a final $2 \times Id^C$ -coalgebra.

Finality is also the key categorical concept underlying the important coalgebraic principle of *coinduction*. Here is a first example.

Example 1.12 Take the function **zip** that merges two streams by taking elements from either stream in turn. For a coalgebraic definition of this map, define the transition map $\delta : (C^\omega \times C^\omega) \rightarrow C \times (C^\omega \times C^\omega)$ as follows:

$$\delta(\alpha, \beta) := (\mathbf{h}(\alpha), (\beta, \mathbf{t}(\alpha))),$$

where $\mathbf{h}$ and $\mathbf{t}$ are the maps defined in Example 1.10. This defines a $K_C \times Id$ -coalgebra on the set $(C^\omega \times C^\omega)$, so that by finality of the stream coalgebra (C^ω, γ) there is a (unique) map $\mathbf{zip} : C^\omega \times C^\omega \rightarrow C^\omega$ which is a coalgebra morphism from $(C^\omega \times C^\omega, \delta)$ to $(C^\omega, \langle \mathbf{h}, \mathbf{t} \rangle)$:

$$\begin{array}{ccc} C^\omega \times C^\omega & \xrightarrow{\mathbf{zip}} & C^\omega \\ \downarrow \delta & & \downarrow \langle \mathbf{h}, \mathbf{t} \rangle \\ C \times (C^\omega \times C^\omega) & \xrightarrow{(\text{id}_C, \mathbf{zip})} & C \times C^\omega \end{array}$$

One may verify that this coalgebra morphism indeed defines the map that zips two streams together.

Unfortunately, as we will see further on, final coalgebras do not exist for every functor.

Example 1.13 In the categories of Kripke frames and Kripke models, final objects do not exist. The *canonical model* comes close, but to turn this structure into a final coalgebra, we have to enrich the base category **Set** with topological structure. As a result that we will discuss later on, we may see the canonical *general* frame as a final coalgebra for a suitable base category and coalgebra functor.

1.4 Behavioural equivalence and bisimilarity

Probably the most intuitive notion of equivalence between systems is that of *behavioral*, or *observational*, equivalence. The idea here is to consider two states to be similar if we cannot distinguish them by observations, because they display the same behavior. For instance, we call two deterministic automata (pointed $2 \times Id^C$ -coalgebras) equivalent if they recognize the same language. In case the functor T admits a final coalgebra $\mathbb{Z}$, this idea is easily formalized by making state s in coalgebra $\mathbb{S}$ equivalent to state s' in coalgebra $\mathbb{S}'$ if $\mathbf{beh}_{\mathbb{S}}(s) = \mathbf{beh}_{\mathbb{S}'}(s')$. In case the functor does not admit a final coalgebra, we generalize this demand as follows.

Definition 1.14 Let $\mathbb{S} = (S, \sigma)$ and $\mathbb{S}' = (S', \sigma')$ be two systems for the set functor T . Then $s \in S$ and $s' \in S'$ are *behaviorally equivalent*, notation: $\mathbb{S}, s \simeq_T \mathbb{S}', s'$ if there is a T -system $\mathbb{X} = (X, \xi)$ and homomorphisms $f : \mathbb{S} \rightarrow \mathbb{X}$ and $f' : \mathbb{S}' \rightarrow \mathbb{X}$ such that $f(s) = f'(s')$. $\triangleleft$

Remark 1.15 It is easily checked that in case T admits a final coalgebra $\mathbb{Z}$, then indeed $\mathbb{S}, s \simeq_T \mathbb{S}', s'$ iff $\mathbf{beh}_{\mathbb{S}}(s) = \mathbf{beh}_{\mathbb{S}'}(s')$. The direction from right to left is trivial, so assume that s and s' are behaviorally equivalent because of the existence of $\mathbb{X}$, f and f' as in the formulation of the definition. Observe that the map $\mathbf{beh}_{\mathbb{X}} \circ f$ is a coalgebra morphism from $\mathbb{S}$ to $\mathbb{Z}$, and likewise for $\mathbf{beh}_{\mathbb{X}} \circ f'$ and $\mathbb{S}'$. It then follows from the finality of $\mathbb{Z}$ that $\mathbf{beh}_{\mathbb{S}} = \mathbf{beh}_{\mathbb{X}} \circ f$ and $\mathbf{beh}_{\mathbb{S}'} = \mathbf{beh}_{\mathbb{X}} \circ f'$. Hence, from $f(s) = f'(s')$ it follows that $\mathbf{beh}_{\mathbb{S}}(s) = \mathbf{beh}_{\mathbb{S}'}(s')$, as required.

As we will see further on, in many cases of interest, behavioral equivalence can be characterized via the equally fundamental concept of *bisimilarity*, which involves the notion of a *coalgebraic bisimulation*.

Definition 1.16 A *bisimulation* between two systems $(\mathbb{S}, \sigma)$ and $(\mathbb{S}', \sigma')$ is a relation $B \subseteq S \times S'$ for which, as in the diagram below,

$$\begin{array}{ccccc}
 S & \xleftarrow{\pi} & B & \xrightarrow{\pi'} & S' \\
 \sigma \downarrow & & \beta \downarrow & & \downarrow \sigma' \\
 TS & \xleftarrow{T\pi} & TB & \xrightarrow{T\pi'} & TS'
 \end{array} \tag{2}$$

there is a (not necessarily unique) coalgebra map $\beta : B \rightarrow TB$ such that the two projection maps, π and π' , from B to S and S' respectively, are both coalgebra morphisms.

We write $\mathbb{S}, s \Leftrightarrow \mathbb{S}', s'$ to denote that there is some bisimulation B between $\mathbb{S}$ and $\mathbb{S}'$ that links the states s in $\mathbb{S}$ and s' in $\mathbb{S}'$. $\triangleleft$

That is, a bisimulation is a relation that, seen as a set, can itself be endowed with a coalgebra structure satisfying some natural conditions.

1.5 Coalgebra and modal logic

Logic comes in when we want to design specification languages for describing the behaviour of state-based evolving systems, and derivation systems for reasoning about this behaviour.

Definition 1.17 An (*abstract*) *logic* for T -coalgebras is a pair $(L, \Vdash)$ consisting of a set L of *formulas* and, for each T -coalgebra $\mathbb{S} = (S, \sigma)$, a *satisfaction* relation $\Vdash_{\mathbb{S}} \subseteq S \times L$. In case $(s, \varphi) \in \Vdash_{\mathbb{S}}$, we say that φ is *true* or *holds* at $s \in \mathbb{S}$, or that s *satisfies* φ in $\mathbb{S}$; we often write $\mathbb{S}, s \Vdash \varphi$ or even $s \Vdash \varphi$ instead of $s \Vdash_{\mathbb{S}} \varphi$.

Given a state s in a coalgebra $\mathbb{S}$, we define $Th_{\mathbb{S}}^L(s) := \{\varphi \in L \mid \mathbb{S}, s \Vdash \varphi\}$. Conversely, given a formula $\varphi \in L$ and a coalgebra $\mathbb{S}$, we let $\llbracket \varphi \rrbracket^{\mathbb{S}}$ denote the set of states in $\mathbb{S}$ where φ holds, that is, $\llbracket \varphi \rrbracket^{\mathbb{S}} := \{s \in S \mid \mathbb{S}, s \Vdash \varphi\}$. If $Th_{\mathbb{S}}^L(s) = Th_{\mathbb{S}'}^L(s')$, we say that s and s' are *L-equivalent*, and we write $\mathbb{S}, s \equiv_L \mathbb{S}', s'$ (or simply $s \equiv_L s'$ if $\mathbb{S}$ and $\mathbb{S}'$ are understood).

Finally, we will call a formula *satisfiable* if it is satisfied at some state in some coalgebra, and *valid* if it holds at every state in every coalgebra. $\triangleleft$

In the same way that universal coalgebra tries to give an account of state-based evolving systems, *uniformly* in the coalgebra type T , research in coalgebraic logic has been directed towards a development of logical languages and derivation systems that are similarly uniform in the parameter T . Apart from uniformity, here are some other desiderata for a coalgebraic logic.

Definition 1.18 Let $(L, \Vdash)$ be a logic for a coalgebra type T . We say that this logic is *invariant* (for behavioural equivalence) if $\mathbb{S}, s \simeq \mathbb{S}', s'$ implies $\mathbb{S}, s \equiv_L \mathbb{S}', s'$ and *expressive* if conversely, $\mathbb{S}, s \equiv_L \mathbb{S}', s'$ implies that s and s' are behaviourally equivalent. The logic is *decidable* if there is an algorithm that decides, on input $\varphi \in L$, whether there is some pointed coalgebra satisfying φ . $\triangleleft$

In addition, for practical purposes one generally wants the logic to be *finitary* in the sense that formulas are finite objects. Other desirable properties of a coalgebraic logic include good model-theoretic behaviour, and the existence of a derivation system that is sound and complete for the collection of valid formulas.

With Kripke models as paradigmatic examples of coalgebra, and modal logic providing *the* bisimulation-invariant logic for Kripke models, it should come as no surprise that most coalgebraic logics can be seen as generalisations of basic modal logic in some sense. The literature on coalgebra witnesses different ways to generalise basic modal logic from Kripke structures to arbitrary systems; here we mention two approaches.

First, however, we briefly discuss the role of *proposition letters* in coalgebraic modal logic. Generalising the relation between Kripke models and Kripke frames, we introduce the notion of a T -model.

Definition 1.19 Let T be a set functor, and let $\mathbf{Q}$ be an arbitrary but fixed set of proposition letters. A T -model is a triple (S, σ, V) such that (S, σ) is a T -coalgebra, and $V : \mathbf{Q} \rightarrow PS$ is a *valuation*. A morphism between T -models $\mathbb{S} = (S, \sigma, V)$ and $\mathbb{S}' = (S', \sigma', V')$ is a coalgebra morphism $f : (S, \sigma) \rightarrow (S', \sigma')$ such that $s \in V(p)$ iff $fs \in V'(p)$, for all $s \in S$. $\triangleleft$

Remark 1.20 There are two natural ways to think about T -models: either as T -coalgebras extended with a $\mathbf{Q}$ -valuation, or as coalgebras for the functor $T_{\mathbf{Q}} := K_{P\mathbf{Q}} \times T$. (Clearly, in the latter case it would be more natural to represent the valuation V as its associated *colouring* $V^b : S \rightarrow P\mathbf{Q}$ given by $V^b(s) := \{p \in \mathbf{Q} \mid s \in V(p)\}$, cf. Example A.5.) In these notes we will generally take the first perspective, since it is more compatible with the perspective on proposition letters as *variables*. Nevertheless we will apply various coalgebraic definitions to T -models as if they were indeed $K_{P\mathbf{Q}} \times T$ -coalgebras.

Let us now take a quick look at two of the approaches towards coalgebraic modal logic.

Example 1.21 In the first approach towards coalgebraic modal logic, which is completely parametric in the functor T , the set of formulas L is closed under the following clause, which introduces a modal operator ∇ :

if $\alpha \in TX$ for some finite set X of formulas, then $\nabla\alpha$ is a formula.

In the case of the powerset functor ($T = P$), we can write, for instance, $\nabla\{\varphi_0, \varphi_1\}$, where φ_0 and φ_1 are formulas. The formula $\nabla\{\varphi_0, \varphi_1\}$ will be equivalent to $(\Diamond\varphi_0 \wedge \Diamond\varphi_1) \wedge \Box(\varphi_0 \vee \varphi_1)$. In general, the semantics of ∇ in a Kripke structure $\mathbb{S} = (S, \sigma)$ will be given as

$$\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } (\sigma(s), \alpha) \in \overline{P}(\Vdash),$$

where $\bar{P}(\Vdash) \subseteq PS \times PL$ is the (Egli-Milner) *lifting*² of the binary satisfaction relation $\Vdash$.

As we will see, this approach generalises well to any set functor T that ‘preserves weak pullbacks’ — the point of this condition being that T preserves weak pullbacks iff its lifting $\bar{T}$ preserves relation composition.

The ∇ -logic provided by the relation-lifting approach described in Example 1.21 may provide coalgebraic logics in a completely uniform way, but its unusual syntax makes it not easy to work with.

The second approach to coalgebraic logic provides coalgebraic logics with a more standard modal syntax. Here, the modalities of the language correspond to so-called *predicate liftings*, where an n -ary predicate lifting is a natural transformation $\check{P}^n \rightarrow \check{P}T$. Here we confine ourselves to a few examples of such coalgebraic modalities.

Example 1.22 The standard interpretation of the modalities $\Diamond$ and $\Box$ in Kripke structures can be formulated as follows:

$$\begin{aligned} \mathbb{S}, s \Vdash \Box\varphi & \text{ iff } R(s) \subseteq \llbracket \varphi \rrbracket \\ \mathbb{S}, s \Vdash \Diamond\varphi & \text{ iff } R(s) \cap \llbracket \varphi \rrbracket \neq \emptyset. \end{aligned}$$

Example 1.23 *Monotone modal logic* is a variant of standard modal logic where formulas are interpreted in so-called monotone neighbourhood models. These are structures of the form $\mathbb{S} = (S, \sigma, V)$ where S is a set of states, V is a valuation, and σ is a map $S \rightarrow PPS$ that assigns to each state $s \in S$ a collection $\sigma(s) \subseteq PS$ of neighbourhoods. Here, each collection $\sigma(s)$ is required to be *upwards closed* in the sense that $X \in \sigma(s)$ implies $Y \in \sigma(s)$ for all Y with $X \subseteq Y \subseteq S$.

In these structures we may interpret the modalities $\Diamond$ and $\Box$ as follows:

$$\begin{aligned} \mathbb{S}, s \Vdash \Box\varphi & \text{ iff } \llbracket \varphi \rrbracket \in \sigma(s) \\ \mathbb{S}, s \Vdash \Diamond\varphi & \text{ iff } (S \setminus \llbracket \varphi \rrbracket) \notin \sigma(s). \end{aligned}$$

Using the upwards-closedness of $\sigma(s) \subseteq P(S)$ it is not hard to show that $\Box\varphi$ holds at s iff s has a neighbourhood $U \in \sigma(s)$ such that $\mathbb{S}, u \Vdash \varphi$, for each $u \in U$, whereas $\Diamond\varphi$ holds at s iff every neighbourhood $U \in \sigma(s)$ contains some point u where φ holds.

To see how monotone modal logic generalises standard modal logic, think of a Kripke model (S, R, V) as the neighbourhood model $(S, \hat{\sigma}, V)$, where $\hat{\sigma}(s) := \{X \in PS \mid R(s) \subseteq X\}$.

Example 1.24 Let (S, σ) be a Markov chain, that is, a coalgebra for the distribution functor D . Given a rational number $q \in [0, 1]$, we introduce a modality $\Diamond_q$, with the following intended meaning:

$$\mathbb{S}, s \Vdash \Diamond_q\varphi \text{ iff } \sum_{t \in \llbracket \varphi \rrbracket^{\mathbb{S}}} \mu_s(t) > q.$$

That is, the formula $\Diamond_q\varphi$ holds at s iff the probability that φ holds at the next state after s is bigger than q .

²The Egli-Milner lifting of a relation $R \subseteq S \times S'$ is the relation $\bar{P}(R) \subseteq PS \times PS'$ given by $(X, X') \in \bar{P}(R)$ iff for all $x \in X$ there is an $x' \in X'$ such that Rxx' and for all $x' \in X'$ there is an $x \in X$ such that Rxx' .

1.6 Literature

Here are some relevant texts on coalgebra and modal logic. First we mention some books:

- J. Barwise and L. Moss, *Vicious Circles*, CSLI Publications, 1996.
- B. Jacobs, *Introduction to Coalgebra: towards mathematics of states and observation*, Cambridge University Press, 2016.
- J. Rutten, *The Method of Coalgebra: exercises in coinduction*, CWI, Amsterdam, The Netherlands, 2019, ISBN 978-90-6196-568-8.

Here is a list of introductory and survey articles:

- B. Jacobs and J. Rutten, *A tutorial on (co)algebras and (co)induction*, Bulletin of the European Association for Theoretical Computer Science, 62 (1997), pp. 222–259..
- J.J.M.M Rutten, *Universal coalgebra: a theory of systems*, Theoretical Computer Science 249 (2000), pp. 3-80.
- H.P. Gumm, *Universelle coalgebra*, appendix to Th. Ihringer: *Universelle Algebra*, Heldermann Verlag, Berlin, 2003. (available from the web page of H.P. Gumm.)
- A. Kurz, A. Palmigiano and Y. Venema, *Coalgebra and logic: an overview*, Journal of Logic and Computation 20 (2010), pp. 985-988.
- C. Kupke and D. Pattinson, *Coalgebraic Semantics of Modal Logics: an Overview*, Theoretical Computer Science, 412(38) (2011), pp. 5070-5094.
- C. Cîrstea, A. Kurz, D. Pattinson, L. Schröder and Y. Venema, *Modal logics are coalgebraic*, The Computer Journal, 54 (2011), pp. 31-41.

2 Final Coalgebras and Coinduction

In section 1.3 we introduced final coalgebras. In this chapter we study the concept in more detail, and we see how it relates to the fundamental coalgebraic definition and proof principle of *coinduction*.

As a first example of a final coalgebra, it is instructive to look at a base category different from **Set**.

Example 2.1 Let $\mathbb{C} = (C, \leq)$ be an arbitrary poset, that is, $\leq$ is a reflexive, transitive and antisymmetric relation on the set C . We may think of $\mathbb{C}$ as a category by taking C as the set of objects and providing a unique arrow between any pair of elements $c, d \in C$ for which $c \leq d$. An *endofunctor* on $\mathbb{C}$ is then nothing but a *monotone* or *order-preserving* function $F : C \rightarrow C$. Given that arrows between objects are unique if they exist, a *coalgebra* $(c, \gamma : c \rightarrow Fc)$ for such a functor can be identified with its carrier c , and conversely, any $c \in C$ for which $c \leq Fc$ is the carrier of an F -coalgebra. In other words, we may identify F -coalgebras with the *prefixpoints* of F .

We leave it for the reader to verify that a final coalgebra for a functor F on $\mathbb{C}$ is a *greatest fixpoint* of the map F .

2.1 The language coalgebra

As a key example of a final coalgebra we will show how to endow the collection of *languages* over some finite alphabet C with coalgebra structure that turns it into the final coalgebra for the set functor $2 \times Id^C$ associated with deterministic automata.

Here we will represent a deterministic automaton over an alphabet C as a triple $\mathbb{S} = (S, \chi, \tau)$, where $\chi : S \rightarrow 2$ and $\tau : S \rightarrow S^C$ correspond to the acceptance condition and the transition map, respectively. Note that we drop the condition that the carrier of the automaton is finite. We will also use the convention that $s \xrightarrow{a} t$ means $t = \tau(s)(a)$ and $s \downarrow$ indicates that s is accepting, i.e., $\chi(s) = 1$.

As we saw in the introduction, we can identify deterministic automata with coalgebras of the functor $2 \times Id^C$. It is easy to see that a map $f : S \rightarrow S'$ is a coalgebra morphism between two automata $\mathbb{S} = (S, \chi, \tau)$ and $\mathbb{S}' = (S', \chi', \tau')$ if it satisfies, for all $s \in S$ and $a \in C$, the conditions $\chi(s) = \chi'(fs)$ and $f(\tau(s)(a)) = \tau'(fs)(a)$.

Definition 2.2 Consider the following *language coalgebra* $\mathbb{L} := (\mathcal{L}_C, \omega, \delta)$, where

- $\mathcal{L}_C := P(C^*)$ is the collection of all *languages* over C ,
- $\omega : \mathcal{L}_C \rightarrow 2$ is given by $\omega(L) := 1$ if $\varepsilon \in L$, and $\omega(L) = 0$ if $\varepsilon \notin L$;
- $\delta : \mathcal{L}_C \rightarrow (\mathcal{L}_C)^C$ is the map given by $\delta(L)(a) := L_a$, the so-called *a-derivative* of L :

$$L_a := \{u \in C^* \mid au \in L\}.$$

If no confusion concerning the alphabet is likely, we will usually write $\mathcal{L}$ rather than $\mathcal{L}_C$. $\triangleleft$

Recall that for an arbitrary automaton $\mathbb{S} = (S, \chi, \tau)$, we defined the language recognized by a state $s \in S$ by putting

$$L_{\mathbb{S}}(s) := \{u \in C^* \mid \chi(\hat{\tau}(s)(u)) = 1\},$$

where $\hat{\tau} : S \rightarrow S^{C^*}$ is inductively defined by putting $\hat{\tau}(s)(\varepsilon) := s$ and $\hat{\tau}(s)(cu) := \hat{\tau}(\tau(s)(c))(u)$.

We claim that, for any alphabet C , $\mathbb{L}$ is the final coalgebra of type $2 \times Id^C$, with the language maps as the witnessing coalgebra morphisms.

Proposition 2.3 (Finality of $\mathbb{L}$) *For any $2 \times Id^C$ -coalgebra $\mathbb{S}$, the map $L_{\mathbb{S}}$ is the unique coalgebra morphism $L_{\mathbb{S}} : \mathbb{S} \rightarrow \mathbb{L}$.*

Proof. Fix $\mathbb{S} = (S, \chi, \tau)$. We first show that $L_{\mathbb{S}}$ is a coalgebra morphism. For acceptance, we check that $\chi(s) = \omega(L_{\mathbb{S}}(s))$:

$$\begin{aligned} \chi(s) = 1 & \text{ iff } \varepsilon \in L_{\mathbb{S}}(s) & (\text{definition } L_{\mathbb{S}}) \\ & \text{ iff } \omega(L_{\mathbb{S}}(s)) = 1 & (\text{definition } \omega) \end{aligned}$$

With respect to the transition function, we need to show that $L_{\mathbb{S}}(\tau(s)(c)) = \delta(L_{\mathbb{S}}(s))(c)$, for all $s \in S$ and $c \in C$. But this identity holds because of the following chain of equivalences, for an arbitrary word $u \in C^*$:

$$\begin{aligned} u \in L_{\mathbb{S}}(\tau(s)(c)) & \text{ iff } \chi(\hat{\tau}(\tau(s)(c))(u)) = 1 & (\text{definition } L_{\mathbb{S}}) \\ & \text{ iff } \chi(\hat{\tau}(s)(cu)) = 1 & (\text{definition } \hat{\tau}) \\ & \text{ iff } cu \in L_{\mathbb{S}}(s) & (\text{definition } L_{\mathbb{S}}) \\ & \text{ iff } u \in \delta(L_{\mathbb{S}}(s))(c) & (\text{definition } \delta) \end{aligned}$$

Second, we prove uniqueness. Assuming that $f : \mathbb{S} \rightarrow \mathbb{L}$ is a coalgebra morphism, we need to show that $f = L_{\mathbb{S}}$. It suffices to show that any word $u \in C^*$ satisfies the following:

$$\text{for all } s \in S : \quad u \in L_{\mathbb{S}}(s) \text{ iff } u \in f(s). \quad (3)$$

We will prove (3) by induction on u . In the base case, where $u = \varepsilon$, we have

$$\begin{aligned} \varepsilon \in L_{\mathbb{S}}(s) & \text{ iff } \omega(L_{\mathbb{S}}(s)) = 1 & (\text{definition } \omega) \\ & \text{ iff } \chi(s) = 1 & (L_{\mathbb{S}} \text{ is a morphism}) \\ & \text{ iff } \omega(f(s)) = 1 & (f \text{ is a morphism}) \\ & \text{ iff } \varepsilon \in f(s) & (\text{definition } \omega) \end{aligned}$$

Now assume that $u = cv$, for some $c \in C$, then we find

$$\begin{aligned} cv \in L_{\mathbb{S}}(s) & \text{ iff } \chi(\hat{\tau}(s)(cv)) = 1 & (\text{definition } L_{\mathbb{S}}) \\ & \text{ iff } \chi(\hat{\tau}(\tau(s)(c))(v)) = 1 & (\text{definition } \hat{\tau}) \\ & \text{ iff } v \in L_{\mathbb{S}}(\tau(s)(c)) & (\text{definition } L_{\mathbb{S}}) \\ & \text{ iff } v \in f(\tau(s)(c)) & (\text{induction hypothesis}) \\ & \text{ iff } v \in \delta(f(s))(c) & (f \text{ is a morphism}) \\ & \text{ iff } cv \in f(s) & (\text{definition } \delta) \end{aligned}$$

This suffices to prove the induction step of (3).

QED

2.2 Properties of final coalgebras

Final coalgebras have various interesting properties. We first show that, if existing, final coalgebras are unique modulo isomorphism³. Because of this fact we will often speak of ‘the’ final T -coalgebra if T admits final coalgebras.

Proposition 2.4 *Let $\mathbb{Z} = (Z, \zeta)$ and $\mathbb{Z}' = (Z', \zeta')$ be final T -coalgebras for some functor $T : \mathbf{C} \rightarrow \mathbf{C}$. Then $\mathbb{Z}$ and $\mathbb{Z}'$ are isomorphic.*

Proof. By finality of $\mathbb{Z}$ there is a coalgebra morphism $g : \mathbb{Z}' \rightarrow \mathbb{Z}$, and by finality of $\mathbb{Z}'$ there is a coalgebra morphism $f : \mathbb{Z} \rightarrow \mathbb{Z}'$. But then the composition $g \circ f : \mathbb{Z} \rightarrow \mathbb{Z}$ is a coalgebra morphism as well, and by unicity it must be identical to the identity arrow $\text{id}_{\mathbb{Z}}$. Similarly, we find that $g \circ f = \text{id}_{\mathbb{Z}'}$. Thus $\mathbb{Z}$ and $\mathbb{Z}'$ are isomorphic indeed. QED

The following proposition states a key fact about final coalgebras.

Proposition 2.5 (Lambek’s Lemma) *Let $\mathbb{Z}$ be a final T -coalgebra for some functor $T : \mathbf{C} \rightarrow \mathbf{C}$. Then the coalgebra map $\zeta : Z \rightarrow TZ$ of $\mathbb{Z}$ is an isomorphism in $\mathbf{C}$.*

Proof. Applying the functor T to the coalgebra map ζ of $\mathbb{Z}$, we obtain the map $T\zeta : TZ \rightarrow TTZ$, and hence, a coalgebra $\mathbb{Z}_2 := (TZ, T\zeta)$. By finality of $\mathbb{Z}$ we obtain a coalgebra morphism $!$ from $\mathbb{Z}_2$ to $\mathbb{Z}$, given by a $\mathbf{C}$ -arrow $! : TZ \rightarrow Z$. But then the composition $! \circ \zeta$ is a coalgebra morphism from $\mathbb{Z}$ to itself, just like the identity arrow id_Z . In a diagram:

$$\begin{array}{ccccc}
 & & \text{id}_Z & & \\
 & \nearrow \zeta & & \searrow ! & \\
 Z & \xrightarrow{\quad} & TZ & \xrightarrow{\quad} & Z \\
 \zeta \downarrow & & \downarrow T\zeta & & \downarrow \zeta \\
 TZ & \xrightarrow{T\zeta} & TTZ & \xrightarrow{T!} & TZ
 \end{array} \tag{4}$$

It follows by unicity that $! \circ \zeta = \text{id}_Z$.

For the reverse composition $\zeta \circ !$ we have that $\zeta \circ ! = T! \circ T\zeta$ since $!$ is a morphism, cf. the right rectangle in the diagram above. But then we easily derive that $\zeta \circ ! = T(! \circ \zeta) = T\text{id}_Z = \text{id}_{TZ}$. In other words, $\zeta \circ !$ is the identity arrow on TZ .

Finally, since $! \circ \zeta = \text{id}_Z$ and $\zeta \circ ! = \text{id}_{TZ}$ we see that ζ is an isomorphism indeed, with $!$ as its inverse. QED

As an immediate corollary of this, we see that set functors involving the full powerset functor in a nontrivial way, will generally not admit a final coalgebra.

Corollary 2.6 *The categories of Kripke frames and of Kripke models do not admit final coalgebras.*

³See the appendix for the categorical definition of an isomorphism. In the case of coalgebras, two coalgebras $\mathbb{S}$ and $\mathbb{S}'$ are isomorphic if there are coalgebra morphisms $f : \mathbb{S} \rightarrow \mathbb{S}'$ and $g : \mathbb{S}' \rightarrow \mathbb{S}$ such that $g \circ f$ and $f \circ g$ are the identity arrows on $\mathbb{S}$ and $\mathbb{S}'$, respectively.

Proof. Recall that Kripke frames are coalgebras for the powerset functor P . Now suppose for contradiction that P admits a final coalgebra $\mathbb{Z} = (Z, \zeta)$. It would follow by Lambek's Lemma that ζ is a bijection between Z and its powerset; but this is impossible by Cantor's theorem. The case of Kripke models can be proved similarly. QED

The following proposition is one way to formalise the proof principle of *coinduction* — we shall come back to this.

Proposition 2.7 *Let $\mathbb{Z}$ be a final T -coalgebra for some set functor T . Then the relation $\simeq_{\mathbb{Z}}$ of behavioural equivalence on $\mathbb{Z}$ is the identity relation Δ_Z on Z :*

$$\simeq_{\mathbb{Z}} = \Delta_Z. \quad (5)$$

Proof. Suppose that z and z' are two states in $\mathbb{Z}$ that are behaviourally equivalent. In Remark 1.15 we saw that this implies that $\text{beh}_{\mathbb{Z}}(z) = \text{beh}_{\mathbb{Z}}(z')$, and since $\text{beh}_{\mathbb{Z}}$ is the identity map on Z , it follows that $z = \text{beh}_{\mathbb{Z}}(z) = \text{beh}_{\mathbb{Z}}(z') = z'$. QED

2.3 Existence of final coalgebras

If not all set functors admit final coalgebras, which ones do? Some good sufficient conditions are known.

Definition 2.8 Let T be some set functor, and κ some cardinal. Call T κ -small if

$$T(S) = \bigcup \{ (T\iota_S^A)[T(A)] \mid \iota_S^A : A \hookrightarrow S, |A| < \kappa \},$$

for all sets $S \neq \emptyset$, where for an arbitrary subset A of S , the arrow ι_S^A denotes the inclusion map of A into S . T is *small* if it is small for some cardinal κ . An ω -small functor is usually called *finitary*. $\triangleleft$

In words, the definition requires every element of $T(S)$ to be in the range of $T\iota$ for an appropriate inclusion map $\iota : A \hookrightarrow S$, where A is of size smaller than κ . In case T *preserves inclusions* (meaning that T maps any inclusion $\iota : A \hookrightarrow B$ to an inclusion $T\iota : TA \hookrightarrow TB$), the definition boils down to the requirement that

$$T(S) = \bigcup \{ T(A) \mid A \subseteq S, |A| < \kappa \}.$$

Fact 2.9 *Every small set functor admits a final coalgebra.*

Examples of small functors abound; for instance, whenever we replace, in a Kripke polynomial functor, the power set functor by a bounded variant such as the finite power set functor, the result is a small functor.

In particular, the finite power set functor P_ω itself is ω -small. As an immediate corollary of this fact, the categories of image-finite frames and models, which can be represented as coalgebras for, respectively, the functors P_ω and $PQ \times P_\omega$, both have final objects. More in general we can prove the following.

Corollary 2.10 *Every finitary Kripke polynomial functor admits a final coalgebra.*

For **Set**-based functors that do not admit a final coalgebra, there are various ‘second-best’ ways to proceed. For instance, one may show that T *does* have a final coalgebra in an *extended* or *modified* category.

Example 2.11 If one is willing to allow coalgebras taking a *class* rather than a set as their carrier, one may *create* a final coalgebra, outside the category **Set**, as follows. Let T be a set functor for which final coalgebras do not exist; for convenience we assume that all functors preserve inclusions.

Let **SET** be the category that has classes as objects, and class functions as arrows, that is, functions that map sets to sets but may have a class rather than a set as their (co-)domain. Call an endofunctor T on **SET** *set-based* if for each class C and each $X \in TC$ there is a set $S \subseteq C$ such that $X \in TS$. Now Aczel & Mendler proved that every set-based endofunctor on **SET** admits a final coalgebra – the similarity to Fact 2.9 is no coincidence.

This fact can be used as follows. Given an endofunctor T on **Set**, there is a *unique* way to extend T to a set-based endofunctor T^+ on **SET**. (On objects, simply put $T^+(C) := \bigcup \{T(S) \mid S \subseteq C \text{ a set}\}$.)

The theorem of Aczel & Mendler then guarantees the existence of a final object $\mathbb{Z}$ in $\text{Coalg}(T^+)$. This coalgebra will be class-based if T does not admit a final coalgebra, but it will be final, not only with respect to the set-based coalgebras in $\text{Coalg}(T^+)$, but also with respect to the class-based ones. As an important manifestation of this idea, Aczel showed that the class of non-well-founded sets provides the final coalgebra for (the **SET**-based extension of) the power set functor.

Example 2.12 One way to look at Lambek’s Lemma is that final T -coalgebras provide solutions to the ‘equation’ $S \cong TS$. In the case of the powerset functor, Cantor’s theorem states that this equation does not have a solution in the category **Set**. This situation is reminiscent of that in *domain theory*, which provides solutions to the equation $X \cong X^X$ by imposing topological structure on sets.

Something similar can be done here. Define a *Stone space* to be a pair $\mathbb{X} = (X, \tau)$, where τ is a zero-dimensional compact Hausdorff space, and let **Stone** denote the category of Stone spaces as objects with continuous maps as arrows. As an analog to the powerset functor on **Set**, we can define the *Vietoris* functor V on the category **Stone**; on objects, the Vietoris space $V\mathbb{X}$ is based on the collection of *compact* subsets of X . We may then show that, indeed, the final coalgebra for this functor exists.

Further on we will see that in modal logic, whereas the canonical model (over a finite set Q of proposition letters) is not final in the category of coalgebras for the Kripke model functor $K_{PQ} \times P$, we may identify the canonical *general* model over Q with the final $K_{PQ} \times V$ -coalgebra. Thus we can solve the equation $S \cong TS$ by *modifying* the base category of our coalgebras.

2.4 The terminal sequence

Whether the functor admits a final coalgebra or not, one may always (try to) approximate it by considering the so-called final or terminal *sequence*.

Example 2.13 Let us first consider an example outside the category **Set**. Suppose that the poset $\mathbb{C} = (C, \leq)$ is in fact a complete lattice, that is, with each subset $X \subseteq C$ we may associate a *meet* or *greatest lower bound* $\bigwedge X$; this means in particular that $\mathbb{C}$ is bounded: it has a largest element $\top := \bigwedge \emptyset$, and a smallest element $\perp := \bigvee \emptyset$.

The Knaster-Tarski theorem states that in this setting, every monotone map $F : C \rightarrow C$ has both a least and a greatest fixpoint. That is, every endofunctor F on the *category* $\mathbb{C}$ admits both an initial and a final coalgebra (see Example 2.1).

It is instructive for our purposes to prove this theorem, and in particular, to see how to find the greatest fixpoint by approximating it from above. We define an ordinal-indexed sequence $\langle z_\alpha \rangle$ using transfinite induction:

$$z_\alpha := \begin{cases} \top & \text{if } \alpha = 0 \\ Fz_\beta & \text{if } \alpha = \beta + 1 \text{ is a successor ordinal} \\ \bigwedge \{z_\beta \mid \beta < \alpha\} & \text{if } \alpha \text{ is a limit ordinal} \end{cases}$$

Note that in fact, if we take 0 to be a limit ordinal, we can reduce the first clause of the definition to the third.

It is not hard to prove, for a monotone map $F : C \rightarrow C$, the existence of an ordinal α for which $z_\alpha = z_{\alpha+1}$, and to show that the object z_α is in fact the greatest fixpoint of F .

In the case of set functors, we may take inspiration from this example to define the *terminal sequence* associated with T .

Definition 2.14 The *final* or *terminal sequence* associated with a given set functor T , is an ordinal indexed sequence of objects $\langle Z_\alpha \rangle$ with maps $p_\beta^\alpha : Z_\alpha \rightarrow Z_\beta$ for $\beta \leq \alpha$, such that (i) $Z_{\alpha+1} = TZ_\alpha$ and $p_{\beta+1}^{\alpha+1} = Tp_\beta^\alpha$, (ii) $p_\alpha^\alpha = \text{id}_{Z_\alpha}$ and $p_\gamma^\beta \circ p_\beta^\alpha = p_\gamma^\alpha$, (iii) if λ is a limit ordinal, then Z_λ with $\{p_\alpha^\lambda \mid \alpha < \lambda\}$ is a limit⁴ of the diagram with objects $\{Z_\alpha \mid \alpha < \lambda\}$ and arrows $\{p_\beta^\alpha \mid \alpha, \beta < \lambda\}$. (In particular, taking 0 to be a limit ordinal, we find that $Z_0 = 1$ is some final object of the category **Set**, i.e., a singleton set.) $\triangleleft$

In the diagram below we depict an initial part of this construction:

$$\begin{array}{ccccccc} & & & & p_0^\omega & & \\ & & & & \swarrow & \searrow & \\ & & & & p_1^\omega & & \\ & & & & \swarrow & \searrow & \\ & & & & p_2^\omega & & \\ & & & & \swarrow & \searrow & \\ Z_0 = 1 & \xleftarrow{p_0^1} & Z_1 = T1 & \xleftarrow{p_1^2} & Z_2 = T^2 1 & \xleftarrow{\quad} & \dots & \xleftarrow{p_\omega^{\omega+1}} & Z_\omega & \xleftarrow{\quad} & Z_{\omega+1} = TZ_\omega & \xleftarrow{\quad} & \dots \end{array}$$

(6)

It is not hard to prove that, modulo isomorphism, the terminal sequence is uniquely determined by these conditions. Intuitively, it can be seen as an approximation of the final coalgebra for T . That is, where elements of the final coalgebra represent ‘complete’ behavior, elements of Z_α represent behavior that can be ‘performed in α many steps’.

⁴See the appendix for the categorical definition of a limit.

To make this precise and formal, observe that for any T -coalgebra $\mathbb{S}$ there is a *unique* ordinal-indexed class of functions $\text{beh}_\alpha : S \rightarrow Z_\alpha$ such that beh_0 is fixed by the finality of Z_0 in $\mathbf{Set}$, $\text{beh}_{\alpha+1} := (T\text{beh}_\alpha) \circ \sigma$:

$$\begin{array}{ccccc}
 \dots & & Z_\alpha & \xleftarrow{p_\alpha^{\alpha+1}} & Z_{\alpha+1} & & \dots \\
 & \uparrow \text{beh}_\alpha & & & & \uparrow T\text{beh}_\alpha & \\
 & S & \xrightarrow{\sigma} & TS & & & \\
 & & \nearrow \text{beh}_{\alpha+1} & & & &
 \end{array} \tag{7}$$

while for a limit ordinal λ , beh_λ is given as the unique map $\text{beh}_\lambda : S \rightarrow Z_\lambda$ such that $\text{beh}_\alpha = p_\alpha^\lambda \circ \text{beh}_\lambda$ for all $\alpha < \lambda$. It is not hard to prove that, for instance, $\mathbb{S}, s \simeq \mathbb{S}', s'$ implies that $\text{beh}_\alpha(s) = \text{beh}_\alpha(s')$ for all α .

The relation with final coalgebras can be made precise, as follows. On the one hand, if the terminal sequence converges, in the sense that some arrow $p_\alpha^{\alpha+1}$ is a bijection, then the coalgebra $(Z_\alpha, (p_\alpha^{\alpha+1})^{-1})$ is a final coalgebra for T . And conversely, under some mild constraints on T , Adámek & Koubek proved that if T admits a final coalgebra, then the terminal sequence converges to it.

2.5 Coinduction as a definition principle

Coinduction is an important coalgebraic principle, and just like its algebraic counterpart of *induction*, it can be used as a tool to *define* various operations, but also as a coalgebraic *proof* principle.

To see how coinductive definitions work, suppose that $\mathbb{Z} = (Z, \zeta)$ is the final coalgebra for some set functor T . Coinduction is based on the observation that, in order to define a map from some set S to Z , it suffices to turn S into a coalgebra by endowing it with coalgebra structure: any coalgebra map $\sigma : S \rightarrow TS$ canonically induces a map from S to Z , namely the unique coalgebra morphism $!_\sigma : (S, \sigma) \rightarrow \mathbb{Z}$. (In fact, we saw this principle already at work in the proof of Lambek's Lemma, where we defined a map from TZ to Z by considering the coalgebra map $T\zeta$ on TZ .)

Example 2.15 Let $\mathbb{Z} = (C^\omega, \mathbf{h}, \mathbf{t})$ be the *stream coalgebra* of Example 1.10 — here we write $\mathbf{h}$ and $\mathbf{t}$ rather than *head* and *tail*. We already saw that $\mathbb{Z}$ is the final $K_C \times \text{Id}$ -coalgebra, we can now use the finality to define operations on streams.

To start with, consider the coalgebra map $\langle \mathbf{h}, \mathbf{t} \circ \mathbf{t} \rangle : C^\omega \rightarrow C \times C^\omega$ defined as $\gamma_e(\alpha) := (\mathbf{h}(\alpha), \mathbf{t}(\mathbf{t}(\alpha)))$. By finality of $\mathbb{Z}$ there is a unique map $\mathbf{e} : C^\omega \rightarrow C^\omega$ making the following diagram commute:

$$\begin{array}{ccc}
 C^\omega & \xrightarrow{\mathbf{e}} & C^\omega \\
 \downarrow \langle \mathbf{h}, \mathbf{t} \circ \mathbf{t} \rangle & & \downarrow \langle \mathbf{h}, \mathbf{t} \rangle \\
 C \times C^\omega & \xrightarrow{\text{id}_C \times \mathbf{e}} & C \times C^\omega
 \end{array} \tag{8}$$

Another way of looking at this definition is that, in order to define $\mathbf{e}(\alpha)$, we specify

$$\begin{aligned} \mathbf{h}(\mathbf{e}(\alpha)) &:= \mathbf{h}(\alpha) \\ \mathbf{t}(\mathbf{e}(\alpha)) &:= \mathbf{t}(\mathbf{t}(\alpha)). \end{aligned}$$

In fact, the map $\mathbf{e} : C^\omega \rightarrow C^\omega$ is the operation on streams that creates a new stream out of all items at an *even* position in the input stream. To prove this, it suffices to show that the map $\mathbf{e} := \lambda\alpha.(\lambda n.\alpha(2n))$ makes the diagram (8) commute.

Similarly, we can define a map $\mathbf{q} : C^\omega \rightarrow C^\omega$ selecting the *odd* items of an input stream, by means of the following diagram:

$$\begin{array}{ccc} C^\omega & \xrightarrow{\mathbf{q}} & C^\omega \\ \langle \mathbf{hot}, \mathbf{tot} \rangle \downarrow & & \downarrow \langle \mathbf{h}, \mathbf{t} \rangle \\ C \times C^\omega & \xrightarrow{\text{id}_C \times \mathbf{q}} & C \times C^\omega \end{array} \quad (9)$$

Example 2.16 Fix an alphabet C . Recall from Proposition 2.3 that the language coalgebra $\mathbb{L} = (\mathcal{L}, \omega, \delta)$ is the final coalgebra for the ‘automaton’ functor $2 \times \text{Id}^C$. We can use this fact to define operations on languages.

For instance, given a word $u = c_1 \cdots c_k$ (with $k \geq 0$), we let $\otimes u$ denote its converse, $\otimes u := c_k \cdots c_1$, and we set $\otimes L := \{\otimes u \mid u \in L\}$. Coinductively, we can define this language by imposing the following structure on $\mathcal{L}$. As the acceptance condition we simply take the same map ω as for $\mathbb{L}$, while for the transition map τ we put

$$\tau(L)(c) := \{u \in C^* \mid uc \in L\}.$$

We leave it as an exercise for the reader to verify that the following diagram commutes:

$$\begin{array}{ccc} \mathcal{L} & \xrightarrow{\otimes} & \mathcal{L} \\ \langle \omega, \tau \rangle \downarrow & & \downarrow \langle \omega, \delta \rangle \\ 2 \times \mathcal{L}^C & \xrightarrow{\text{id}_2 \times \otimes^C} & 2 \times \mathcal{L}^C \end{array} \quad (10)$$

so that by finality, $\otimes : \mathcal{L} \rightarrow \mathcal{L}$ is the *unique* coalgebra morphism $\otimes : (\mathcal{L}, \tau, \omega) \rightarrow \mathbb{L}$. That means that (10) can be seen as a coinductive *definition* of $\otimes$.

Example 2.17 As a slightly different example, we give a coinductive definition of the *shuffle* $K \parallel L$ of two languages K and L . For an intuitive description of this operation⁵, a word belongs to $K \parallel L$ if it consists of two interleaved words, from K and L , respectively. In particular, we want that

$$(K \parallel L)_c = K_c \parallel L \cup K \parallel L_c,$$

⁵For a more precise definition of the language $K \parallel L$, first define the relation $\triangleleft$ on finite words by putting $c_1 \cdots c_k \triangleleft d_1 \cdots d_n$ if there is an order-preserving map on the indices $f : \{1, \dots, k\} \rightarrow \{1, \dots, n\}$ such that $c_i = d_{f_i}$, for all $i \in \{1, \dots, k\}$. Say that w *merges* u and v if $|w| = |u| + |v|$ and both $u \triangleleft w$ and $v \triangleleft w$. We can then define $K \parallel L$ as the collection of all words that merge words from K and L .

for all $c \in C$, where we recall that, for a language $L \in \mathcal{L}$, we let L_a denote its a -derivative, $L_a := \{u \in C^* \mid au \in L\}$.

For a coinductive definition, we consider the set $\mathcal{E}$ of expressions defined by the following grammar:

$$E ::= \underline{L} \mid E_0 + E_1 \mid E_0 \parallel E_1$$

where $L \in \mathcal{L}$, i.e., we associate a formal symbol $\underline{L}$ with every language L .

To turn the set $\mathcal{E}$ into an automaton $\mathbb{E} := (\mathcal{E}, \chi, \tau)$, consider the following axioms and rules, which inductively define the relation $\downarrow$ and the functions $\xrightarrow{c}$, with $c \in C$. Here we use the notation introduced in the beginning of section 2.1.

$$\begin{array}{c} \frac{}{\underline{L} \downarrow} \varepsilon \in L \quad \frac{E \downarrow}{(E + F) \downarrow} \quad \frac{F \downarrow}{(E + F) \downarrow} \quad \frac{E \downarrow \quad F \downarrow}{(E \parallel F) \downarrow} \\[10pt] \frac{}{\underline{L} \xrightarrow{c} \underline{L}_c} \quad \frac{E \xrightarrow{c} E' \quad F \xrightarrow{c} F'}{E + F \xrightarrow{c} E' + F'} \quad \frac{E \xrightarrow{c} E' \quad F \xrightarrow{c} F'}{E \parallel F \xrightarrow{c} E' \parallel F + E \parallel F'} \end{array}$$

It is not hard to see that this deductive system uniquely determines two operations, $\chi : \mathcal{E} \rightarrow 2$ and $\tau : \mathcal{E} \rightarrow \mathcal{E}^C$. Hence by finality of $\mathbb{L}$ there is a unique coalgebra morphism $f : \mathbb{E} \rightarrow \mathbb{L}$. Think of $f(E)$ as the *interpretation* of the term E . We can then define

$$K \parallel L := f(\underline{K} \parallel \underline{L}).$$

To get a feeling for this definition we compute the derivative $(K \parallel L)_c$:

$$\begin{aligned} (K \parallel L)_c &= f(\underline{K} \parallel \underline{L})_c && \text{(definition } \parallel) \\ &= \delta(f(\underline{K} \parallel \underline{L}))(c) && \text{(definition } \delta) \\ &= f(\tau(\underline{K} \parallel \underline{L}))(c) && (f \text{ is a morphism}) \\ &= f(\underline{K}_c \parallel \underline{L} + \underline{K} \parallel \underline{L}_c) && \text{(definition } \tau) \\ &= f(\underline{K}_c \parallel \underline{L}) \cup f(\underline{K} \parallel \underline{L}_c) && (*) \\ &= K_c \parallel L \cup K \parallel L_c && \text{(definition } \parallel) \end{aligned}$$

Here we use in $(*)$ the observation that $f(E + F) = f(E) \cup f(F)$, which is easily proved by coinduction, see Example 2.20.

2.6 Coinduction as a proof principle

As a *proof* principle, coinduction has two manifestations. In its most direct form, proofs by coinduction use the uniqueness of coalgebra morphisms to a final coalgebra.

Example 2.18 Consider the maps e and q of Example 2.15. We claim that

$$q = e \circ t, \tag{11}$$

as should be clear intuitively. To prove (11) coinductively, it suffices to prove that the map $e \circ t : C^\omega \rightarrow C^\omega$, just like q , is a coalgebra morphism $e \circ t : (C^\omega, \langle h \circ t, t \circ t \rangle) \rightarrow (C^\omega, \langle h, t \rangle)$; that is, the diagram below commutes:

$$\begin{array}{ccc} C^\omega & \xrightarrow{e \circ t} & C^\omega \\ \langle h \circ t, t \circ t \rangle \downarrow & & \downarrow \langle h, t \rangle \\ C \times C^\omega & \xrightarrow{id_C \times e \circ t} & C \times C^\omega \end{array} \quad (12)$$

But for this purpose it suffices to show that the following two equations hold:

$$h \circ (e \circ t) = h \circ t \quad (13)$$

$$t \circ (e \circ t) = (e \circ t) \circ (t \circ t) \quad (14)$$

This is not so hard. For (13), we may calculate

$$\begin{aligned} h \circ (e \circ t) &= (h \circ e) \circ t && \text{(associativity)} \\ &= (id_C \circ h) \circ t && \text{(diagram (8))} \\ &= h \circ t, && (id_C \circ h = h) \end{aligned}$$

while we prove (14) as follows:

$$\begin{aligned} t \circ (e \circ t) &= (t \circ e) \circ t && \text{(associativity)} \\ &= (e \circ (t \circ t)) \circ t && \text{(diagram (8))} \\ &= (e \circ t) \circ (t \circ t) && \text{(associativity)} \end{aligned}$$

Example 2.19 In a similar way we can prove that

$$zip \circ \langle e, q \rangle = id_{C^\omega}, \quad (15)$$

where zip is the map defined in Example 1.12.

By finality it suffices to show that $zip \circ \langle e, q \rangle$ is a coalgebra morphism on the stream coalgebra, and since we know that zip is a coalgebra morphism (i.e., the right rectangle below commutes), we can confine ourselves to proving that the left rectangle in the diagram below commutes:

$$\begin{array}{ccccc} C^\omega & \xrightarrow{\langle e, q \rangle} & C^\omega \times C^\omega & \xrightarrow{zip} & C^\omega \\ \langle h, t \rangle \downarrow & & \downarrow \delta & & \downarrow \langle h, t \rangle \\ C \times C^\omega & \xrightarrow{T\langle e, q \rangle} & C \times (C^\omega \times C^\omega) & \xrightarrow{Tzip} & C \times C^\omega \end{array}$$

Here $T\langle e, q \rangle = id_C \times \langle e, q \rangle$, $Tzip = id_C \times zip$, and δ is as given in Example 1.12: $\delta(\alpha, \beta) := (h(\alpha), (\beta, t(\alpha)))$.

To verify that the left rectangle above commutes we need to check that $\delta \circ \langle e, q \rangle = T\langle e, q \rangle \circ \langle h, t \rangle$, which boils down to proving

$$(a) \quad h \circ e = h,$$

- (b) $q = e \circ t$
- (c) $t \circ e = q \circ t$.

But we obtain (a) because $h \circ e = \text{id}_C \circ h$ (definition of e), and (b) was shown in the previous example, cf. (11). Finally, for (c), observe that $t \circ e = e \circ (t \circ t)$ by definition of e (diagram (8)), and $e \circ (t \circ t) = q \circ t$ by associativity and (11).

Often, coinduction is referred to as the proof principle that uses the fact that behavioural equivalence is the identity relation on a final coalgebra (Proposition 2.7). More specifically, given the fact that bisimilarity implies behavioural equivalence (as we will see further on), one may prove two states in a final coalgebra to be identical if we can link them by a bisimulation.

Example 2.20 In the case of C -automata (coalgebras of type $2 \times \text{Id}^C$), a bisimulation on a coalgebra (S, χ, τ) is a relation $B \subseteq S \times S$ such that, whenever $(s_0, s_1) \in B$, we have

- (acc) $\chi(s_0) = \chi(s_1)$ (that is: $s_0 \downarrow$ iff $s_1 \downarrow$), and
- (nxt) $(\tau(s_0)(c), \tau(s_1)(c)) \in B$, for all $c \in C$.

Let us now prove the statement that (cf. Example 2.17)

$$f(\underline{L}) = L, \tag{16}$$

for every language $L \in \mathcal{L}$. By coinduction, it suffices to show that the relation

$$B := \{(f(\underline{L}), L) \mid L \in \mathcal{L}\} \tag{17}$$

is a bisimulation. We check the two conditions.

For (acc), we observe that $f(\underline{L}) \downarrow$ (in $\mathbb{L}$) iff $\underline{L} \downarrow$ (in $\mathbb{E}$) since f is a coalgebra morphism. But we have $\underline{L} \downarrow$ in $\mathbb{E}$ iff $\varepsilon \in L$ by definition of acceptance in $\mathbb{E}$, and we have $L \downarrow$ (in $\mathbb{L}$) iff $\varepsilon \in L$ by definition of acceptance in $\mathbb{L}$. This suffices to prove (acc).

For (nxt) we need to show, for an arbitrary language $L \in \mathcal{L}$ and an arbitrary letter $c \in C$, that the pair $(\delta(f(\underline{L}), c), \delta(L, c)) \in B$. To that aim, observe that $\delta(f(\underline{L}), c) = f(\tau(\underline{L})(c)) = f(\underline{L}_c)$, respectively since f is a morphism and by definition of τ . But since we have $\delta(L)(c) = L_c$ (by definition of δ), it is immediate that $(\delta(f(\underline{L}), c), \delta(L, c)) = (f(\underline{L}_c), L_c) \in B$. This finishes the proofs of (17) and (16).

Similarly, we can prove that

$$f(E_0 + E_1) = f(E_0) \cup f(E_1), \tag{18}$$

for all expressions E_0 and E_1 , by showing that the relation

$$R := \{(f(E_0 + E_1), f(E_0) \cup f(E_1)) \mid E_0, E_1 \in \mathcal{E}\}$$

is a bisimulation on $\mathbb{L}$.

Finally, we can use coinduction to establish, in a relatively straightforward way, various useful properties of the operation $\parallel$, such as commutativity, associativity, or distribution with respect to $+$ / $\cup$.

3 Bisimilarity and Behavioural Equivalence

In section 1.4 of the Introduction we defined two coalgebraic notions of equivalence: behavioral equivalence and bisimilarity. In this chapter we discuss these notions in more detail.

3.1 Basic observations

Obviously, the first question is how the notions of behavioral equivalence and bisimilarity relate to each other. One direction is clear: bisimilarity is a sufficient condition for behavioral equivalence.

Proposition 3.1 *Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be some functor, and let s_0 and s_1 be states of the T -coalgebras $\mathbb{S}_0$ and $\mathbb{S}_1$, respectively. Then $\mathbb{S}_0, s_0 \Leftrightarrow \mathbb{S}_1, s_1$ implies $\mathbb{S}_0, s_0 \simeq \mathbb{S}_1, s_1$.*

Proof. In the special case that T admits a final coalgebra, a very simple proof obtains. Assume that $\mathbb{S}_0, s_0 \Leftrightarrow \mathbb{S}_1, s_1$, and let $B \subseteq S_0 \times S_1$ with $\beta : B \rightarrow TB$ be a coalgebra witnessing this. It follows from the definitions that both $\text{beh}_{\mathbb{S}_0} \circ \pi_0$ and $\text{beh}_{\mathbb{S}_1} \circ \pi_1$ are coalgebraic morphisms from (B, β) to the final coalgebra, so from finality it follows that $\text{beh}_{\mathbb{S}_0} \circ \pi_0 = \text{beh}_{\mathbb{S}_1} \circ \pi_1$. From this it is immediate that $B \subseteq \simeq$; and so from $(s_0, s_1) \in B$ it follows that $\mathbb{S}_0, s_0 \simeq \mathbb{S}_1, s_1$.

In the general case the proof of this proposition is similar to the one of Theorem 3.13 below (with an application of *pushouts* instead of pullbacks), so we omit details. QED

The converse statement of Proposition 3.1 does not hold: in general, bisimilarity is a strictly stronger notion than behavioral equivalence.

Example 3.2 Consider the so-called ‘3-2-functor’ $T_2^3 : \mathbf{Set} \rightarrow \mathbf{Set}$ given on objects by

$$T_2^3(S) := \{(s_0, s_1, s_2) \in S^3 \mid |\{s_0, s_1, s_2\}| \leq 2\},$$

while for an arrow $f : S \rightarrow S'$ we define $(T_2^3 f)(s_0, s_1, s_2) := (fs_0, fs_1, fs_2)$. We leave it as an exercise to the reader to verify that this indeed defines a set functor.

Now consider the following coalgebra $\mathbb{S} = (S, \sigma)$, where $S = \{0, 1\}$ and σ is given by $\sigma(0) = (0, 0, 1)$ and $\sigma(1) = (1, 0, 0)$. Then it is not hard to see that $\mathbb{S}, 0 \simeq \mathbb{S}, 1$, but at the same time we claim that there is *no* T_2^3 -bisimulation on $\mathbb{S}$ linking 0 and 1. To see this, suppose for contradiction that $R \subseteq S \times S$ would be such a bisimulation, witnessed by the coalgebra map $\rho : R \rightarrow T_2^3 R$. If the projection maps $\pi_0, \pi_1 : R \rightarrow S$ are to be coalgebra morphisms, ρ has to map the pair $(0, 1)$ to some triple $\rho(0, 1) = ((s_0, s_1), (t_0, t_1), (u_0, u_1))$ such that

$$\begin{aligned} (s_0, t_0, u_0) &= ((T_2^3 \pi_0) \circ \rho)(0, 1) = (\sigma \circ \pi_0)(0, 1) = \sigma(0) = (0, 0, 1) \\ (s_1, t_1, u_1) &= ((T_2^3 \pi_1) \circ \rho)(0, 1) = (\sigma \circ \pi_1)(0, 1) = \sigma(1) = (1, 0, 0). \end{aligned}$$

Clearly then we find $\rho(0, 1) = ((s_0, s_1), (t_0, t_1), (u_0, u_1)) = ((0, 1), (0, 0), (1, 0))$. But this object does not belong to the set $T_2^3 R$, since (s_0, s_1) , (t_0, t_1) and (u_0, u_1) are all distinct.

Example 3.3 A more natural example of a set functor for which behavioural equivalence and bisimilarity are properly distinct notions, is the *monotone neighbourhood functor* M . We will come back to this example later.

In section 3.3 below we will discuss an important class of set functors for which we do have $\simeq = \Leftrightarrow$. First, however, we make some basic observations on bisimulations and, in the next section, we give an alternative characterization of bisimulations.

Example 3.4 For an arbitrary set functor T , it is easy to see that for any coalgebra $\mathbb{S}$, the diagonal relation $\Delta_{\mathbb{S}}$ is a bisimulation equivalence on $\mathbb{S}$. Furthermore, the converse of a bisimulation is again a bisimulation.

As another general example, coalgebra morphisms can be seen as functional bisimulations. To be more precise, let $f : S_0 \rightarrow S_1$ be a function between the carriers of two T -coalgebras $\mathbb{S}_0$ and $\mathbb{S}_1$. Recall that the *graph* of f is the relation $\text{Gr}f := \{(s, f(s)) \mid s \in S_0\}$. Then it holds that

$$f \text{ is a coalgebra morphism iff its graph } \text{Gr}f \text{ is a bisimulation.} \quad (19)$$

In order to see why this is so, first suppose that $\text{Gr}f : \mathbb{S}_0 \Leftrightarrow \mathbb{S}_1$. Since the projection map $\pi_0 : \text{Gr}f \rightarrow S_0$ is a bijective morphism, its inverse π_0^{-1} is also a morphism. But then $f = \pi_1 \circ \pi_0^{-1}$, as the composition of two morphisms, is also a morphism. For the other direction, suppose that f is a morphism; then it is straightforward to verify that the map $(T\pi_0)^{-1} \circ \sigma \circ \pi_0$ equips the set $\text{Gr}f$ with the required coalgebraic structure.

However, the collection of bisimulations is not in general closed under taking relational composition, and the relation $\Leftrightarrow$ of bisimilarity on a given coalgebra is generally not an equivalence relation.

3.2 Bisimulations and relation lifting

Bisimulations admit an elegant alternative characterization which involves the notion of *relation lifting*.

Example 3.5 As an example, consider the power set functor P . Recall that a relation $B \subseteq S_0 \times S_1$ is a bisimulation between two P -coalgebras (Kripke frames) $\mathbb{S}_0 = (S_0, R_0[\cdot])$ and $\mathbb{S}_1 = (S_1, R_1[\cdot])$ iff B satisfies the conditions (*back*) and (*forth*) of Example 1.3. Now suppose that we define, for an arbitrary relation $R \subseteq S_0 \times S_1$, the relation $\bar{P}(R) \subseteq P(S_0) \times P(S_1)$ by putting

$$\bar{P}(R) := \{(Q_0, Q_1) \mid \forall q_0 \in Q_0 \exists q_1 \in Q_1. (q_0, q_1) \in R \text{ and } \forall q_1 \in Q_1 \exists q_0 \in Q_0. (q_0, q_1) \in R\}. \quad (20)$$

In other words, we *lift* the relation R to the level of the power sets of S_0 and S_1 . The definition of a bisimulation between P -coalgebras can now be characterized as follows:

$$B : \mathbb{S}_0 \Leftrightarrow \mathbb{S}_1 \text{ iff } (R_0[s_0], R_1[s_1]) \in \bar{P}(B) \text{ for all } (s_0, s_1) \in B.$$

This nice way of characterizing bisimulation via relation lifting is not limited to the power set functor — it applies in fact to *every* set functor.

Definition 3.6 Let T be some set functor. Given a relation $R \subseteq S_0 \times S_1$, consider R as a *span*

$$S_0 \xleftarrow{\pi_0} R \xrightarrow{\pi_1} S_1 ,$$

where $\pi_i : R \rightarrow S_i$ and $p_i : TS_0 \times TS_1 \rightarrow TS_i$ denote the respective projection maps. We define the *relation lifting* of R as the relation $\overline{TR} \subseteq TS_0 \times TS_1$ given by

$$\overline{TR} := \{((T\pi_0)(u), (T\pi_1)(u)) \mid u \in TR\}, \quad (21)$$

that is, $\overline{TR}$ is the image of TR under the map $\tau_R := \langle T\pi_0, T\pi_1 \rangle$. $\triangleleft$

In other words, we apply the functor T to the relation R , seen as a span. It follows from the category-theoretic properties of the product $TS_0 \times TS_1$ that there is a unique map $\tau_R := \langle T\pi_0, T\pi_1 \rangle$ from TR to $TS_0 \times TS_1$ such that $p_i \circ \tau_R = T\pi_i$ for $i = 0, 1$. Now we define $\overline{TR}$ as the image of TR under the map τ_R obtained from the lifted projection maps $T\pi_0$ and $T\pi_1$. In a diagram:

$$\begin{array}{ccccc} TS_0 & \xleftarrow{T\pi_0} & TR & \xrightarrow{T\pi_1} & TS_1 \\ & \searrow p_0 & \downarrow \tau_R & \swarrow p_1 & \\ & & \overline{TR} & & \\ & & \downarrow & & \\ & & TS_0 \times TS_1 & & \end{array}$$

Remark 3.7 Strictly speaking, the definition of the relation lifting of a given relation depends on how we type it. That is, given sets A, B, A' and B' , and a relation R such that $R \subseteq A \times B$ and $R \subseteq A' \times B'$, for the definition of TR it matters whether we look at R as a relation from A to B or as a relation from A' to B' .

This possible source of ambiguity evaporates if we require the functor T to be standard, see the appendix for more details. In this chapter we will only use the lifting of a relation in a setting where its type is fixed, but readers who are worried about this issue may add the (fairly harmless) condition that T is standard.

The results listed in the following theorem summarize the most important properties of bisimulations.

Theorem 3.8 *Let $\mathbb{S}_0$ and $\mathbb{S}_1$ be two coalgebras for some set functor T .*

1. *For any set $B \subseteq S_0 \times S_1$, we have $B : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1$ iff $(\sigma_0(s_0), \sigma_1(s_1)) \in \overline{T}(B)$, all $(s_0, s_1) \in B$.*
2. *The collection of bisimulations between $\mathbb{S}_0$ and $\mathbb{S}_1$ forms a complete lattice under the inclusion order, with joins given by unions.*
3. *The bisimilarity relation $\rightleftharpoons$ is the largest bisimulation between $\mathbb{S}_0$ and $\mathbb{S}_1$.*

Proof. The first part of the theorem is an almost immediate consequence of the definitions. To see this, recall that we have $B : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1$ iff we can find a coalgebra map $\beta : B \rightarrow TB$ such that $(T\pi_i) \circ \beta = \sigma_i \circ \pi_i$ for $i = 0, 1$, and that the latter requirement is equivalent to stating that $(T\pi_i)(\beta(s_0, s_1)) = \sigma_i s_i$. From this it easily follows that $B : \mathbb{S}_0 \rightleftharpoons \mathbb{S}_1$ iff for every

$(s_0, s_1) \in B$ there is a $u \in TB$ such that $(\sigma_0 s_0, \sigma_1 s_1) = ((T\pi_0)(u), (T\pi_1)(u))$. This suffices by (79).

The crucial observation in the proof of the other two parts is that

$$\bar{T} : P(S_0 \times S_1) \rightarrow P(TS_0 \times TS_1) \text{ is a monotone operation.} \quad (22)$$

For a proof, let $R \subseteq R'$ be two relations between S_0 and S_1 , with $\iota : R \rightarrow R'$ denoting the inclusion map. By definition of $\bar{T}$, we may without loss of generality represent an arbitrary element of $\bar{T}(R)$ as a pair $\tau_R(u) = ((T\pi_0)(u), (T\pi_1)(u))$ for some $u \in TR$. Define $u' := (T\iota)(u)$, then u' belongs to TR' , and for each i we find that $(T\pi'_i)(u') = (T\pi'_i \circ T\iota)(u) = T(\pi'_i \circ \iota)(u) = (T\pi_i)(u)$. That is, $\tau_R(u) = \tau_{R'}(u')$, which shows that $\tau_R(u)$ belongs to $\bar{T}R'$. This proves (22).

Now for the proof of part 2, recall that a partial order is a complete lattice if it is closed under arbitrary joins. Hence, it suffices to prove that the union B of a collection $\{B_j \mid j \in J\}$ of bisimulations is again a bisimulation. Take an arbitrary pair $(s_0, s_1) \in B$. Then (s_0, s_1) belongs to B_j for some $j \in J$. Hence, by part 1, we find $(\sigma_0 s_0, \sigma_1 s_1) \in \bar{T}(B_j)$, so $(\sigma_0(s_0), \sigma_1(s_1)) \in \bar{T}(B)$ by the monotonicity of $\bar{T}$. But then B is a bisimulation by part 1.

Finally, for part 3, note that it is an immediate consequence of part 2 that $\Leftrightarrow$, being the union of all bisimulations between $\mathbb{S}_0$ and $\mathbb{S}_1$, is a bisimulation itself. Hence, by definition, it is the greatest bisimulation between $\mathbb{S}_0$ and $\mathbb{S}_1$. In fact, it follows by the Knaster-Tarski theorem (on fixed points of monotone operations on complete lattices), that $\Leftrightarrow$ is in fact the greatest fixed point of the map $\Lambda : R \mapsto \{(s_0, s_1) \mid (\sigma_0(s_0), \sigma_1(s_1)) \in \bar{T}(R)\}$. QED

In the case of Kripke polynomial functors, relation lifting can be characterized using *induction* on the construction of the functor.

Proposition 3.9 *Let S and S' be two sets, and let $R \subseteq S \times S'$ be a binary relation between S and S' . Then the following induction defines the relation lifting $\bar{K}(R) \subseteq KS \times KS'$, for each Kripke polynomial functor K :*


$$\begin{aligned} \bar{Id}(R) &:= R, \\ \bar{K}_C(R) &:= \Delta_C, \\ \bar{K}_0 \times \bar{K}_1(R) &:= \{((x_0, x_1), (x'_0, x'_1)) \mid (x_0, x'_0) \in \bar{K}_0(R) \text{ and } (x_1, x'_1) \in \bar{K}_1(R)\}, \\ \bar{K}_0 + \bar{K}_1(R) &:= \{(\kappa_0 x_0, \kappa_0 x'_0) \mid (x_0, x'_0) \in \bar{K}_0(R)\} \cup \{(\kappa_1 x_1, \kappa_1 x'_1) \mid (x_1, x'_1) \in \bar{K}_1(R)\}, \\ \bar{K}^D(R) &:= \{(f, f') \mid (f(d), f'(d)) \in \bar{K}(R) \text{ for all } d \in D\}, \\ \bar{PK}(R) &:= \{(Q, Q') \mid \forall q \in Q \exists q' \in Q'. (q, q') \in \bar{K}(R) \text{ and } \forall q' \in Q' \exists q \in Q. (q, q') \in \bar{K}(R)\}. \end{aligned}$$

Here κ_0 and κ_1 are the co-projection maps associated with the coproduct, cf. Definition B.14.

3.3 Bisimilarity and behavioural equivalence: smooth functors

In Example 3.2 we saw that bisimilarity is a strictly stronger notion than behavioural equivalence. Here is a constraint on the functor that guarantees the two notions to coincide.

Definition 3.10 A *weak pullback* of two arrows $f_0 : A_0 \rightarrow B$, $f_1 : A_1 \rightarrow B$ in a category $\mathbf{C}$ is a pair of arrows $p_0 : W \rightarrow A_0$, $p_1 : W \rightarrow A_1$ such that (i) $f_0 \circ p_0 = f_1 \circ p_1$, while (ii) for every pair $p'_0 : W' \rightarrow A_0$, $p'_1 : W' \rightarrow A_1$ that also satisfies $f_0 \circ p'_0 = f_1 \circ p'_1$, there is a mediating arrow $w' : W' \rightarrow W$ such that $p_0 \circ w' = p'_0$ and $p_1 \circ w' = p'_1$. We will call a functor $T : \mathbf{C} \rightarrow \mathbf{C}'$ *smooth* if it preserves weak pullbacks; that is, if for any weak pullback (p_0, p_1) of any (f_0, f_1) in $\mathbf{C}$, the pair (Tp_0, Tp_1) is a weak pullback of (Tf_0, Tf_1) in $\mathbf{C}'$. $\triangleleft$



Note that the mediating arrow w' need not be unique: adding this requirement to the definition would give the more familiar, and stronger, notion of a *pullback*. The category $\mathbf{Set}$ has pullbacks: for $f_0 : A_0 \rightarrow B$ and $f_1 : A_1 \rightarrow B$, we can take the projections to A_0 and A_1 from the set $\mathbf{pb}(f_0, f_1) := \{(a_0, a_1) \in A_0 \times A_1 \mid f_0(a_0) = f_1(a_1)\}$.

Many but not all endofunctors on $\mathbf{Set}$ in fact preserve weak pullbacks.

Proposition 3.11 *All polynomial functors preserve pullbacks, and all Kripke polynomial functors preserve weak pullbacks.*

The main reason that this *prima facie* rather exotic property is in fact of great importance in the theory of universal coalgebra, is the following fact (see also Fact B.37 in Appendix B).

Fact 3.12 *For any set functor T the following are equivalent:*

- (1) T is smooth;
- (2) $\overline{T}(R; Q) = \overline{T}R; \overline{T}Q$, for all pairs of relations $R \subseteq X \times Y$ and $Q \subseteq Y \times Z$.

Theorem 3.13 *If T is a smooth set functor, the following hold on the class of T -coalgebras:*

- (1) *the relational composition of two bisimulations is again a bisimulation;*
- (2) *the notions of bisimilarity and behavioral equivalence coincide:*

$$\simeq_T = \dot{\simeq}_T.$$

Proof. For the proof of the first statement, let $\mathbb{S}_0, \mathbb{S}_1$ and $\mathbb{S}_2$ be T -coalgebras, and let $B' : \mathbb{S}_0 \dot{\simeq} \mathbb{S}_1$ and $B'' : \mathbb{S}_1 \dot{\simeq} \mathbb{S}_2$ be bisimulations. We will show that $B := B'; B''$ is a bisimulation as well, between $\mathbb{S}_0$ and $\mathbb{S}_2$. For that purpose, take an arbitrary pair $(s_0, s_2) \in B$; then by definition there must be a state $s_1 \in \mathbb{S}_1$ such that $(s_0, s_1) \in B'$ and $(s_1, s_2) \in B''$. But since B' and B'' are bisimulations, this means that $(\sigma_0 s_0, \sigma_1 s_1) \in \overline{T}B'$ and $(\sigma_1 s_1, \sigma_2 s_2) \in \overline{T}B''$, so that $(\sigma_0 s_0, \sigma_2 s_2) \in \overline{T}B'; \overline{T}B''$. Hence by smoothness we find that $(\sigma_0 s_0, \sigma_2 s_2) \in \overline{T}(B'; B'') = \overline{T}B$, as required.


Turning to the second statement, let s_0 and s_1 be states of the T -coalgebras $\mathbb{S}_0$ and $\mathbb{S}_1$, respectively. We need to prove that $\mathbb{S}_0, s_0 \dot{\simeq} \mathbb{S}_1, s_1$ iff $\mathbb{S}_0, s_0 \simeq \mathbb{S}_1, s_1$. Because of Proposition 3.1 it suffices to prove the direction from right to left.

Let $f_0 : S_0 \rightarrow X$ and $f_1 : S_1 \rightarrow X$ be two coalgebra morphisms such that $f_0(s_0) = f_1(s_1)$. Then in **Set**, the set $B := \{(s_0, s_1) \in S_0 \times S_1 \mid f_0(s_0) = f_1(s_1)\}$, together with the projection functions $\pi_0 : B \rightarrow S_0$ and $\pi_1 : B \rightarrow S_1$ constitutes a pullback of f_0 and f_1 , cf. the square in the foreground of the picture. Because T preserves weak pullbacks, the square in the background of the picture is then a weak pullback diagram in **Set**.

Now consider the two arrows $\sigma_i \circ \pi_i : B \rightarrow T(S_i)$. First observe that $Tf_i \circ \sigma_i = \xi \circ f_i$ for each i , because each f_i is a coalgebra morphism. Hence, chasing the diagram we find that

$$\begin{aligned} Tf_0 \circ \sigma_0 \circ \pi_0 &= \xi \circ f_0 \circ \pi_0 \\ &= \xi \circ f_1 \circ \pi_1 = Tf_1 \circ \sigma_1 \circ \pi_1. \end{aligned}$$

Since $T\pi_0$ and $T\pi_1$ form a weak pullback of Tf_0 and Tf_1 , this implies the existence of a mediating function $\beta : B \rightarrow TB$ such that $T\pi_i \circ \beta = \sigma_i \circ \pi_i$. In other words, $\mathbb{B} := (B, \beta)$ is a T -coalgebra, and the projection maps π_0 and π_1 are morphisms from $\mathbb{B}$ to S_0 and S_1 , respectively.



QED

Corollary 3.14 *Let T be a smooth functor. Then for any T -coalgebra S , the maximal bisimulation $\Leftrightarrow$ on S is an equivalence relation.*

Proof. It is not hard to see that the maximal bisimulation $\Leftrightarrow$ is always reflexive and symmetric. For transitivity, it follows by smoothness of T that the composition $\Leftrightarrow; \Leftrightarrow$ is again a bisimulation; but since $\Leftrightarrow$ is maximal we find that $\Leftrightarrow; \Leftrightarrow \subseteq \Leftrightarrow$, which is just another formulation of the transitivity of $\Leftrightarrow$.

QED

4 Covarieties

In universal *algebra* an important part is played by varieties: classes of algebras that are closed under the operations of taking homomorphic images, subalgebras and products of algebras. In this chapter, we introduce the notion of a *covariety* as a natural coalgebraic analog of a variety, and we consider some natural closure operations on classes of coalgebras.

4.1 Homomorphic images

Definition 4.1 Let T be some endofunctor on **Set**. If $f : \mathbb{S} \rightarrow \mathbb{S}'$ is a surjective coalgebra morphism between the T -coalgebras $\mathbb{S}$ and $\mathbb{S}'$, then we say that $\mathbb{S}'$ is a *homomorphic image* of $\mathbb{S}$, and write $\mathbb{S} \twoheadrightarrow \mathbb{S}'$, or $f : \mathbb{S} \twoheadrightarrow \mathbb{S}'$ if we want to make the morphism explicit. $\triangleleft$

In universal algebra, one finds a one-one correspondence between homomorphic images and congruences. There is a similar connection here between homomorphic images and bisimulation equivalences, but the analogy is perfect only in the case of functors that preserve weak pullbacks.

Proposition 4.2 Let $\mathbb{S} = (S, \sigma)$ be a T -coalgebra for some set functor T .

- (1) Given a bisimulation equivalence⁶ E on $\mathbb{S}$, there is a unique coalgebra structure $\bar{\sigma}$ on S/E such that the quotient map $q : S \rightarrow S/E$ is a coalgebra morphism.
- (2) If T preserves weak pullbacks, then the relation $\ker(f) := \{(s, t) \in S^2 \mid fs = ft\}$ is a bisimulation equivalence, for any coalgebra morphism $f : \mathbb{S} \rightarrow \mathbb{S}'$.

Proof. For part (1), we leave it as an exercise for the reader to show that the set $\bar{S} = S/E$ of E -cells, together with the quotient map q , is a coequalizer of the projection maps $\pi_0, \pi_1 : E \rightarrow S$:

$$E \begin{array}{c} \xrightarrow{\pi_0} \\ \xrightarrow{\pi_1} \end{array} S \xrightarrow{q} \bar{S}$$

Now assume that, next to being an equivalence relation, E is also a bisimulation on $\mathbb{S}$. Then by definition there is a coalgebra map $\eta : E \rightarrow TE$ such that both π_i are coalgebra morphisms $\pi_i : (E, \eta) \rightarrow \mathbb{S}$. It follows that

$$\begin{aligned} Tq \circ \sigma \circ \pi_0 &= Tq \circ T\pi_0 \circ \eta && (\pi_0 \text{ is a morphism}) \\ &= T(q \circ \pi_0) \circ \eta && (\text{functoriality}) \\ &= T(q \circ \pi_1) \circ \eta && (q \text{ is a coequalizer}) \\ &= Tq \circ T\pi_1 \circ \eta && (\text{functoriality}) \\ &= Tq \circ \sigma \circ \pi_1 && (\pi_1 \text{ is a morphism}) \end{aligned}$$

In other words, the map $Tq \circ \sigma : S \rightarrow T\bar{S}$ is a competitor for the coequalizer map q , and so

⁶A *bisimulation equivalence* is a bisimulation that is also an equivalence relation.

there is a unique map $\bar{\sigma} : \bar{S} \rightarrow T\bar{S}$ such that $\bar{\sigma} \circ q = Tq \circ \sigma$, in a diagram:

$$\begin{array}{ccccc} E & \xrightarrow[\pi_1]{\pi_0} & S & \xrightarrow{q} & \bar{S} \\ \eta \downarrow & & \downarrow \sigma & & \downarrow \bar{\sigma} \\ TE & \xrightarrow[T\pi_1]{T\pi_0} & TS & \xrightarrow{Tq} & T\bar{S} \end{array}$$

Clearly then $\bar{\sigma}$ is the required coalgebra map on $\bar{S}$.

For the second part of the proposition, observe that $\ker(f)$ is the relational composition of the graph of f with its converse. The result then follows from Theorem 3.13. QED

4.2 Subcoalgebras

The next class operation that we consider is that of taking subcoalgebras.

Definition 4.3 Let $\mathbb{X} = (X, \xi)$ be some T -coalgebra. We say that a subset S of X is *open* with respect to $\mathbb{X}$ if there is a coalgebra map $\sigma : S \rightarrow TS$ such that the inclusion map $\iota : S \rightarrow X$ is a coalgebra morphism from (S, σ) to (X, ξ) ; in this case we call the structure (S, σ) a *subcoalgebra* of $\mathbb{X}$, writing $\mathbb{S} \leq \mathbb{X}$. $\triangleleft$

Interestingly enough, the transition map of a subcoalgebra is completely determined by the underlying open set.

Proposition 4.4 Let $\mathbb{S}_0 = (S, \sigma_0)$ and $\mathbb{S}_1 = (S, \sigma_1)$ be two subcoalgebras of the coalgebra $\mathbb{X}$. Then $\sigma_0 = \sigma_1$.

Proof. The case of S being empty is trivial, so suppose otherwise. Then from the assumption that $\mathbb{S}_0$ and $\mathbb{S}_1$ are subcoalgebras of $\mathbb{A}$, we may infer that $(T\iota) \circ \sigma_0 = \xi \circ \iota = (T\iota) \circ \sigma_1$, where ι is the inclusion map of S into X . It follows from the functoriality of T that $T\iota$ is an injection, so that we may conclude that $\sigma_0 = \sigma_1$. QED

Remark 4.5 In case T is a standard functor, we may simplify Definition 4.3: in this case any subcoalgebra of $\mathbb{X} = (X, \xi)$ must be of the form $\mathbb{S} = (S, \xi \upharpoonright_S)$ with $S \subseteq X$ and $\xi(s) \in TS$, for all $s \in S$.

Example 4.6 (a) Given a Kripke frame (S, R) we call a structure (S', R') a *subframe* of $\mathbb{S}$ if $S' \subseteq S$ and $R' = R \cap (S' \times S')$. Such a subframe is *generated* if its universe S' is closed under the relation R , that is, if $s \in S'$ and $(s, t) \in R$ imply $t \in S'$. It is easy to see that the generated subframes of a frame correspond to the subcoalgebras of its coalgebraic representation.

(b) Similarly, a state b in a deterministic finite automaton $\mathbb{A} = (A, \delta, F)$ is *reachable* from a state $a \in A$ if there is some word u such that $\hat{\delta}(a, u) = b$. Given some automaton $\mathbb{A}$, the states that are reachable from a given state a form an open set.

(c) The empty set is always open: every coalgebra has the empty coalgebra as a subcoalgebra.

Some further observations concerning subcoalgebras are in order. First of all, the topological terminology is justified by the following proposition.

Proposition 4.7 *Given a coalgebra $\mathbb{X}$ for some set functor T , the collection $\tau_{\mathbb{X}}$ of $\mathbb{X}$ -open sets forms a topology.*

Proof. Closure of $\tau_{\mathbb{X}}$ under taking (arbitrary) unions follows from Theorem 3.8, together with the observation that


$$S \subseteq X \text{ is open with respect to } \mathbb{X} \text{ iff } \Delta_S \text{ is a bisimulation on } \mathbb{X}, \quad (23)$$

which in its turn is an immediate consequence of (19).

To prove that $\tau_{\mathbb{X}}$ is closed under taking finite intersections, assume that $\mathbb{A} = (A, \alpha)$ and $\mathbb{B} = (B, \beta)$ are two subcoalgebras of $\mathbb{S}$. The case where $A \cap B = \emptyset$ is trivial, so assume otherwise. Define $C := A \cap B$, fix some element $c \in C$, and consider the maps $f_c : A \rightarrow C$ and $g_c : S \rightarrow B$, given by

$$f_c(a) := \begin{cases} a & \text{if } a \in C \\ c & \text{otherwise} \end{cases} \quad \text{and} \quad g_c(s) := \begin{cases} s & \text{if } s \in B \\ c & \text{otherwise} \end{cases}.$$

Now consider the following diagram:



We define the map $\gamma : C \rightarrow TC$ by putting

$$\gamma := Tf_c \circ \alpha \circ \iota_A^C,$$

and claim that the resulting structure (C, γ) is a subcoalgebra of $\mathbb{S} = (S, \sigma)$. For this purpose we have to show that

$$\sigma \circ \iota_S^C = T\iota_S^C \circ \gamma. \quad (24)$$

We leave it as an exercise for the reader to show (24), using equalities such as $\iota_B^C \circ f_c = g_c \circ \iota_S^A$ and $g_c \circ \iota_S^B = \text{id}_B$. QED

Remark 4.8 In the case where T is standard, the proof that $\tau_{\mathbb{X}}$ is closed under taking finite intersections can again be simplified. In order to prove that $C := A \cap B$ is the carrier of a subcoalgebra of $\mathbb{S}$, it suffices to show that $\sigma(c) \in TC$ for an arbitrary element c of $C := A \cap B$. But since T is standard we have $TC = TA \cap TB$, and we find $\sigma(c) \in TA \cap TB$ by the assumptions on A and B , respectively.

It follows from Proposition 4.7 that the partial order $(\tau_{\mathbb{X}}, \subseteq)$, consisting of the collection of open subsets of X ordered by set inclusion actually forms a *complete lattice*. In this lattice, the *join* of a collection $\mathcal{U}$ of opens is simply given as its union, but its *meet*, that is, the largest open set contained in each open $U \in \mathcal{U}$, is not necessarily its intersection.

In particular, given a subset S of X , there is an open set $V \subseteq X$ which is the *meet* of the collection $\{U \in \tau_{\mathbb{X}} \mid S \subseteq U\}$. However, there is no guarantee that V is the *intersection* of this collection, or, indeed, that S is actually a subset of V , as is witnessed by the example below. This example thus shows that it does not make sense to speak of the smallest subcoalgebra containing a given subset, or of *point-generated* subcoalgebras.

It follows from Proposition 4.7 that, given a subset S of (the carrier of) a coalgebra $\mathbb{X}$, there is a largest subcoalgebra $\langle S \rangle_{\mathbb{X}}$ of $\mathbb{X}$ (of which the carrier is) contained in S : Its carrier is given as the union of all open subsets of S . Note however, that this carrier might be empty, even if S is not.

Example 4.9 Consider the standard Euclidean topology on the real numbers, seen as a coalgebra for the *filter functor* F . This functor is a subfunctor of the (monotone) neighborhood functor which maps a set S to the collection of all *filters* on S^7 and a function $f : S \rightarrow S'$ simply to the function $Mf = \check{P}\check{P}f$. Prime examples of F -coalgebras are the topological spaces. To see this, represent a topology σ on the set S by the function mapping a point $s \in S$ to the collection $\{U \in \sigma \mid s \in U\}$ of its neighborhoods.

One can show that a set S of reals is open in the topological sense iff it is open in the sense of Definition 4.3 — in fact, this holds for any topology. Now take an arbitrary point r in $\mathbb{R}$. Obviously, the *meet* of all open neighborhoods containing r is the empty set. This example thus shows that in general it does not make sense to speak of point-generated subcoalgebras.

Before we turn to further coalgebraic constructions, consider the following natural link between homomorphic images and subcoalgebras.

Proposition 4.10 *Given a coalgebra morphism $f : \mathbb{S} \rightarrow \mathbb{S}'$, there is a (unique) subcoalgebra $f[\mathbb{S}]$ of $\mathbb{S}'$ such that $f : \mathbb{S} \rightarrow f[\mathbb{S}]$ is a surjective morphism.*

Proof. For a proof of this proposition, let $X := f[S]$ be the (set-theoretic) image of S under f , and let $g : X \rightarrow S$ be a right inverse of f , that is, $f(g(x)) = x$ for all $x \in X$. Now define $\xi : X \rightarrow TX$ by $\xi := Tf \circ \sigma \circ g$. It can be shown that the resulting structure $\mathbb{X}$ is always a subcoalgebra of $\mathbb{S}'$, and that $f : \mathbb{S} \rightarrow \mathbb{X}$ is a surjective morphism; further details are left for the reader. QED

⁷Recall that a filter on S is a collection $\mathcal{F}$ of subsets of S which is not only upward closed (with respect to the inclusion relation), but also closed under taking (finite) intersections, that is, $X \cap Y \in \mathcal{F}$ if $X, Y \in \mathcal{F}$.

4.3 Sums and other colimits

Our last example of a coalgebraic construction concerns the straightforward generalization of the disjoint union of Kripke models and frames. The idea is embodied in the following Proposition.

Proposition 4.11 *Let $\mathbb{S}_0 = (S_0, \sigma_0)$ and $\mathbb{S}_1 = (S_1, \sigma_1)$ be two T -coalgebras, and let $S := S_0 \uplus S_1$ be the disjoint union of S_0 and S_1 . Then there is a unique arrow $\sigma : S \rightarrow TS$ making the embeddings κ_i into coalgebra morphisms.*

Proof. Consider the diagram below, where S , together with the embedding maps κ_0 and κ_1 , is the coproduct of S_0 and S_1 . Since the maps $T\kappa_0 \circ \sigma_0$ and $T\kappa_1 \circ \sigma_1$ provide an alternative co-cone, there must be a (unique) mediating arrow $\sigma : S \rightarrow TS$, making the two rectangles in the diagram commute.

$$\begin{array}{ccccc}
 S_0 & \xrightarrow{\kappa_0} & S & \xleftarrow{\kappa_1} & S_1 \\
 \sigma_0 \downarrow & & \downarrow \sigma & & \downarrow \sigma_1 \\
 TS_0 & \xrightarrow{T\kappa_0} & TS & \xleftarrow{T\kappa_1} & TS_1
 \end{array} \tag{25}$$

Clearly then this σ meets the requirements stated in the Proposition. QED

The sum of an arbitrary collection of coalgebras is defined as follows.

Definition 4.12 The *sum* $\coprod_I \mathbb{S}_i$ of a family $\{\mathbb{S}_i \mid i \in I\}$ of coalgebras for some set functor T , is defined by endowing the disjoint union $S := \biguplus_I S_i$ with the unique map $\sigma : S \rightarrow TS$ which turns all embeddings $\kappa_i : S_i \rightarrow S$ into coalgebra morphisms. $\triangleleft$

Sums are coproducts in the category of coalgebras.

Proposition 4.13 *Let $\{\mathbb{S}_i \mid i \in I\}$ be a family of coalgebras for some set functor T . Then $\coprod_I \mathbb{S}_i$ is a coproduct of $\{\mathbb{S}_i \mid i \in I\}$ in the category $\mathbf{Coalg}(T)$.*

In particular, Proposition 4.13 implies that coproducts actually *exist* (in the category of coalgebras for a given set functor T).

Remark 4.14 This result can be in fact be generalised to *arbitrary* colimits: it can be shown that for every set functor T , the category $\mathbf{Coalg}(T)$ has all colimits, and that these colimits are in fact based on the underlying colimits in $\mathbf{Set}$.

4.4 Covarieties

We have now gathered all the basic class operations needed to define the notion of a covariety.

Definition 4.15 Let T be some endofunctor on $\mathbf{Set}$. A class of T -coalgebras is a *covariety* if it is closed under taking homomorphic images, subcoalgebras and sums. The smallest covariety containing a class $\mathbf{K}$ of T -coalgebras is called the covariety *generated* by $\mathbf{K}$, notation: $\mathbf{Covar}(\mathbf{K})$. $\triangleleft$

As in the case of universal algebra, in order to obtain a more succinct characterization of the covariety generated by a class of coalgebras, one may develop a calculus of class operations.

Definition 4.16 Let H , S and Σ denote the class operations of taking (isomorphic copies of) homomorphic images, subcoalgebras, and sums, respectively. $\triangleleft$

On the basis of these (and other) operations one may investigate the validity of ‘inequalities’ like $HS \leq SH$ (meaning that $HS(K) \subseteq SH(K)$ for all classes K of coalgebras). We first prove that each of the three operations H , S and Σ is a *closure operation*; in each case the proof of this observation is straightforward.

Proposition 4.17 *Let C be one of the class operations H , S or Σ , and let K and K' be arbitrary classes of T -coalgebras. Then we find*

- 1) $K \subseteq C(K)$;
- 2) $K \subseteq K'$ implies $C(K) \subseteq C(K')$;
- 3) $C(C(K)) \subseteq C(K)$.

In the following proposition we list some relevant inequalities involving class operations.

Proposition 4.18 *Let K be an arbitrary class of T -coalgebras. Then we find*

- 1) $HS(K) \subseteq SH(K)$;
- 2) $\Sigma S(K) = S\Sigma(K)$;
- 3) $\Sigma H(K) \subseteq H\Sigma(K)$.

Proof. For part (1), let $\mathbb{A}$ be a coalgebra in $HS(K)$. Then there is a coalgebra $\mathbb{C}$ in K , and a coalgebra $\mathbb{B}$, with morphisms e and f such that $e : \mathbb{B} \rightarrow \mathbb{C}$ and $f : \mathbb{B} \rightarrow \mathbb{A}$. Now consider the following push-out diagram in $\text{Coalg}(T)$ (cf. Example B.20):

$$\begin{array}{ccc} \mathbb{B} & \xrightarrow{f} & \mathbb{A} \\ e \downarrow & & \downarrow p_A \\ \mathbb{C} & \xrightarrow{p_C} & \mathbb{P} \end{array}$$

Since (P, p_A, p_C) is also the push-out of (B, e, f) in Set , it follows that the maps p_A and p_C are, respectively, surjective and injective.⁸ But then we immediately find that $\mathbb{P} \in H(K)$ and $\mathbb{A} \in SH(K)$, as required.

We leave the proof of part (2) as an exercise for the reader. For part (3), assume that we have $\mathbb{B} \in \Sigma H(K)$. That is, there is a family of algebras $\{\mathbb{A}_i \mid i \in I\}$ in K , and a family of surjective morphisms $\{f_i : \mathbb{A}_i \rightarrow \mathbb{B}_i \mid i \in I\}$ such that $\mathbb{B} = \Sigma_i \mathbb{A}_i$, with the injections $\kappa_i : \mathbb{B}_i \rightarrow \mathbb{B}$.

Let $\mathbb{A}$ be the sum of the coalgebras $\mathbb{A}_i$, with the morphisms $\lambda_i : \mathbb{A}_i \rightarrow \mathbb{A}$. Then the coalgebra $\mathbb{B}$, with the morphisms $\{\kappa_i \circ f_i : \mathbb{A}_i \rightarrow \mathbb{B}_i\}$, is a competitor cone of $\mathbb{A}$, and so there

⁸Some justification is needed here.

is a unique morphism $f : \mathbb{A} \rightarrow \mathbb{B}$ such that $f \circ \lambda_i = \kappa_i \circ f_i$, for all i :

$$\begin{array}{ccc} \mathbb{A}_i & \xrightarrow{f_i} & \mathbb{B} \\ \lambda_i \downarrow & & \downarrow \kappa_i \\ \mathbb{A} & \xrightarrow{f} & \mathbb{B} \end{array}$$

It remains to show that f is surjective. For that purpose, take an arbitrary state $b \in \mathbb{B}$, then b must be of the form $b = \kappa_i(b_i)$ for some $b_i \in B_i$. By surjectivity of f_i there is a state $a \in A_i$ such that $b_i = f_i a_i$. But then we find that $b = \kappa_i f_i a_i = f \lambda_i a_i$, which obviously implies that b lies in the range of f . So f is surjective indeed. QED

The above results easily lead to the following coalgebraic analog of Tarski's HSP-theorem in universal algebra.

Theorem 4.19 *Let $\mathbf{K}$ be a class of T -coalgebras for some set functor T . Then*

$$\text{Covar}(\mathbf{K}) = \text{SH}\Sigma(\mathbf{K}).$$

Proof. It is an immediate consequence of the definitions that $\text{SH}\Sigma(\mathbf{K}) \subseteq \text{Covar}(\mathbf{K})$. For the opposite inclusion it suffices to show that the class $\text{SH}\Sigma(\mathbf{K})$ is closed under the operations $\mathbf{S}$, $\mathbf{H}$ and $\Sigma(\mathbf{K})$, but this is rather obvious by the Propositions 4.17 and 4.18. For instance, we easily verify that $\Sigma(\text{SH}\Sigma(\mathbf{K})) = \mathbf{S}\Sigma\mathbf{H}\Sigma(\mathbf{K}) \subseteq \text{SH}\Sigma\Sigma(\mathbf{K}) = \text{SH}\Sigma(\mathbf{K})$. QED

5 Coalgebraic modalities via relation lifting

In this chapter we take an approach to coalgebraic logic which is completely uniform in the type functor T . We introduce a coalgebraic modality ∇ of which the ‘arity’ is the finitary version T_ω of the functor itself. That is, the set L of formulas will be closed under the following clause:

if $\alpha \in TX$ for some finite set X of formulas, then $\nabla\alpha$ is a formula.

whereas the semantics of ∇ will be defined by *lifting* the satisfaction relation $\Vdash$ between states and formulas to the relation $\overline{T}(\Vdash)$.

In the special case where T is the powerset functor P , the nabla operator ∇ is known under the name of the *cover modality*; we discuss this case in some detail before moving on to the more general case.

Convention 5.1 Throughout this chapter we will assume that T is a smooth and standard set functor; that is, T preserves both weak pullbacks and inclusions. The first restriction is to ensure optimal behaviour of the relation lifting $\overline{T}$, while the second one is mainly for convenience. In the Facts B.36 and B.38 we list a number of properties of the operation $\overline{T}$ (all of which will be used throughout this chapter).

Furthermore, we will assume that $\mathbf{Q}$ is an arbitrary but fixed set of proposition letters.

5.1 The cover modality

As we will see now, there is an interesting coalgebraic alternative for the standard formulation of basic modal logic in terms of boxes and diamonds. This alternative set-up is based on a connective ∇ , sometimes referred to as the *cover modality*, which turns a (*finite*) set α of formulas into a formula $\nabla\alpha$.

Definition 5.2 Formulas of the language $\text{ML}_{\nabla}(\mathbf{Q})$ are given by the following recursive definition:

$$a ::= p \mid \perp \mid \neg a \mid a_0 \vee a_1 \mid \nabla\alpha$$

where $p \in \mathbf{Q}$, and α denotes a finite set of formulas. $\triangleleft$

Observe that formulas will be denoted by lower case letters $a, b, \dots$

For the semantics of the cover modality, observe that we may think of the forcing or satisfaction relation $\Vdash$ simply as a binary relation between states and formulas. This relation can thus be lifted to a relation $\overline{P}(\Vdash)$ between *sets* of formulas and *sets* of states.

Definition 5.3 The semantics of this modality in a Kripke model $\mathbb{S} = (S, R, V)$ is given by

$$\mathbb{S}, s \Vdash \nabla\alpha \text{ iff } (R(s), \alpha) \in \overline{P}(\Vdash),$$

where $\overline{P}(\Vdash)$ denotes the Egli-Milner relation lifting of the relation $\Vdash$. $\triangleleft$

In words: $\nabla\alpha$ holds at s iff every successor of s satisfies some formula in α , and every formula in α holds at some successor of s . The modality ∇ is sometimes called the *cover modality*: it holds at a state s if the set $\{\llbracket a \rrbracket^{\mathbb{S}} \mid a \in \alpha\}$ covers the collection $R(s)$ of successors of s , in the sense that $R(s) \subseteq \bigcup\{\llbracket a \rrbracket^{\mathbb{S}} \mid a \in \alpha\}$, while at the same time $R(s) \cap \llbracket a \rrbracket^{\mathbb{S}} \neq \emptyset$, for every $a \in \alpha$.

Remark 5.4 It is not so hard to see that the cover modality can be defined in the standard modal language:

$$\nabla\alpha \equiv \Box \bigvee \alpha \wedge \bigwedge \Diamond\alpha, \quad (26)$$

where $\Diamond\alpha$ denotes the set $\{\Diamond a \mid a \in \alpha\}$. Things start to get interesting once we realize that both the ordinary diamond $\Diamond$ and the ordinary box $\Box$ can be expressed in terms of the cover modality (and the disjunction):

$$\begin{aligned} \Diamond a &\equiv \nabla\{a, \top\}, \\ \Box a &\equiv \nabla\emptyset \vee \nabla\{a\}. \end{aligned} \quad (27)$$

Here, as always, we use the convention that $\bigvee \emptyset = \perp$ and $\bigwedge \emptyset = \top$.

Given that ∇ and $\{\Diamond, \Box\}$ are mutually expressible, we arrive at the following proposition. Here we say that two languages are *effectively equi-expressive* if there are effectively definable truth-preserving translations from one language to the other, and vice versa. Recall that ML is the language of standard modal logic.

Proposition 5.5 *The languages ML and ML_{∇} are effectively equi-expressive.*

A remarkable observation about the cover modality is that we can do far better than this: based on the following *modal distributive law*, we can almost completely eliminate the Boolean connective of conjunction from the language ML_{∇} .

Proposition 5.6 *Let α and α' be two sets of formulas. Then the following two formulas are equivalent:*

$$\nabla\alpha \wedge \nabla\alpha' \equiv \bigvee_{Z \in \alpha \bowtie \alpha'} \nabla\{a \wedge a' \mid (a, a') \in Z\}, \quad (28)$$

where $\alpha \bowtie \alpha'$ is the set of all binary relations $Z \subseteq \alpha \times \alpha'$ such that $(\alpha, \alpha') \in \overline{P}(Z)$.

Proof. For the direction from left to right, suppose that $\mathbb{S}, s \Vdash \nabla\alpha \wedge \nabla\alpha'$. Let $Z \subseteq \alpha \times \alpha'$ consist of those pairs (a, a') such that the conjunction $a \wedge a'$ is true at some successor t of s . It is then straightforward to verify that Z is full on α and α' , that is: $(\alpha, \alpha') \in \overline{P}(Z)$, and that $\mathbb{S}, s \Vdash \nabla\{a \wedge a' \mid (a, a') \in Z\}$.

The converse direction is a fairly direct consequence of the definitions. QED

As a corollary of Proposition 5.6 we can restrict the use of conjunction in modal logic to that of a *special conjunction* connective $\bullet$ which may only be applied to a propositional formula and a ∇ -formula.

Definition 5.7 We first define the set $\text{CL}(\mathbf{Q})$ of *literal conjunctions* by the following grammar:

$$\pi ::= p \mid \neg p \mid \perp \mid \top \mid \pi \wedge \pi,$$

and then let the following grammar define the set $\text{DML}_{\nabla}(\mathbf{Q})$ of *disjunctive modal formulas* in $\mathbf{Q}$:

$$a ::= p \mid \neg p \mid \perp \mid \top \mid a_0 \vee a_1 \mid \pi \bullet \nabla \alpha.$$

Here $p \in \mathbf{Q}$, $\pi \in \text{CL}(\mathbf{Q})$ and $\alpha \in P_{\omega}\text{DML}_{\nabla}(\mathbf{Q})$. $\triangleleft$

As mentioned, the bullet connective is semantically equivalent to conjunction:

$$\mathbb{S}, s \Vdash \pi \bullet \nabla \alpha \text{ iff } \mathbb{S}, s \Vdash \pi \text{ and } \mathbb{S}, s \Vdash \nabla \alpha.$$

Note however, that this conjunction is special in the sense that it combines ‘local’ information about s itself with information about the unfolding of s . The point is that this is the *only* form of conjunction that the language DML_{∇} allows.

Theorem 5.8 *The languages ML and DML_{∇} are effectively equi-expressive.*

Proof. We will show how to rewrite a formula $a \in \text{ML}$ into an equivalent formula in DML_{∇} . Start by rewriting a into *negation normal form*:

$$a ::= p \mid \neg p \mid \perp \mid \top \mid a_0 \vee a_1 \mid a_0 \wedge a_1 \mid \Diamond a \mid \Box a,$$

then by (27) we can find an equivalent formula a' in the language given by

$$a ::= p \mid \neg p \mid \perp \mid \top \mid a_0 \vee a_1 \mid a_0 \wedge a_1 \mid \nabla \alpha.$$

Finally, the modal distributive law (28) allows us to push down nabla’s to the propositional level, and so using the propositional distributive law $(a \wedge (b_0 \vee b_1)) \equiv (a \wedge b_0) \vee (a \wedge b_1)$, we can rewrite a' into an equivalent disjunctive modal formula. QED

Theorem 5.8 can be used to prove various interesting results about modal logic, such as the finite model property, or the decidability of the satisfiability problem — in *linear* time, once the formula is in disjunctive normal form. Rather than proving these corollaries here, we will prove these results in the far more general setting of the *coalgebraic* cover modality.

5.2 Moss’ coalgebraic cover modality

We will now generalise the cover modality from the case where $T = P$ to the setting where T is an arbitrary smooth and standard functor. We are eager to keep our language *finitary*, in the sense that formulas will be finitary objects, with for instance finitely many subformulas. For this reason we will work with the *finitary version* of the functor T .

Recall from Convention 5.1 that T preserves inclusions; because of this we may define its *finitary version* $T_{\omega} : \mathbf{Set} \rightarrow \mathbf{Set}$ by putting

$$\begin{aligned} T_{\omega}(S) &:= \bigcup \{TX \mid X \subseteq_{\omega} S\}, \\ T_{\omega}(f : S \rightarrow S') &:= (Tf) \upharpoonright_{T_{\omega}S}. \end{aligned}$$

It is easy to verify that T_ω also preserves inclusions; given the definition of T_ω on functions, we may write Tf instead of $T_\omega f$ without causing confusion. Note, however, that T_ω being finitary does not necessarily mean that $T_\omega S$ is *finite*, even if S is so; a simple counterexample is given by the bag functor.

Definition 5.9 Formulas of the language $\text{ML}_T(\mathbf{Q})$ are given by the following recursive definition:

$$a ::= p \mid \perp \mid \neg a \mid a_0 \vee a_1 \mid \nabla \alpha$$

where $p \in \mathbf{Q}$, and $\alpha \in T_\omega(\text{ML}_T(\mathbf{Q}))$. We will often write ML_T instead of $\text{ML}_T(\mathbf{Q})$ if the set $\mathbf{Q}$ of proposition letters is either understood or irrelevant. $\triangleleft$

The semantics of ML_T is defined as follows. Recall that a *T-model* is a triple (S, σ, V) , where (S, σ) is a *T-coalgebra* and $V : \mathbf{Q} \rightarrow PS$ is a valuation.

Definition 5.10 Let $\mathbb{S} = (S, \sigma, V)$ be a *T-model*. Then by induction on the complexity of ML_T -formulas we define the satisfaction relation $\Vdash$:

$$\begin{array}{ll} \mathbb{S}, s \Vdash p & \text{iff } s \in V(p) \\ \mathbb{S}, s \Vdash \perp & : \text{ never} \\ \mathbb{S}, s \Vdash \neg a & \text{iff } \mathbb{S}, s \not\Vdash a \\ \mathbb{S}, s \Vdash a_0 \vee a_1 & \text{iff } \mathbb{S}, s \Vdash a_0 \text{ or } \mathbb{S}, s \Vdash a_1 \\ \mathbb{S}, s \Vdash \nabla \alpha & \text{iff } (\sigma(s), \alpha) \in \overline{T}(\Vdash). \end{array}$$

A formula $a \in \text{ML}_T$ is *satisfiable* iff a is satisfiable in some state of some *T-model*, and *valid* if its negation is not satisfiable.

Furthermore, we say that two pointed *T-models* are ML_T -*equivalent* or (*modally*) *equivalent* if they satisfy the same ML_T -formulas, notation: $\mathbb{S}, s \equiv_T \mathbb{S}', s'$. $\triangleleft$

Remark 5.11 Before we consider the instantiations of this logic for some set functor T , we argue that the semantics of ML_T is *well defined*. The reader might have some worries about the inductive clause for the ∇ modality, since the definition refers to the lifting of the *full* satisfaction relation.

The point is that because of our assumptions on T , its associated relation lifting $\overline{T}$ commutes with restrictions, cf. Fact B.38. This means that

$$(\sigma(s), \alpha) \in \overline{T}(\Vdash) \text{ iff } (\sigma(s), \alpha) \in \overline{T}(\Vdash|_{S \times X}), \quad (29)$$

where X is any finite set of formulas such that $\alpha \in T_\omega X$. Thus, in order to determine whether $\nabla \alpha$ holds at s or not, we only have to know the interpretation of the formulas used in the justification that $\nabla \alpha$ is a formula. Below we shall see that in fact there is a unique set $\text{Base}(\alpha)$ which is the smallest (finite) set X such that $\alpha \in TX$. In other words, we may replace the ‘quasi-inductive’ clause for ∇ in Definition 5.10 with the following, properly inductive one:

$$\mathbb{S}, s \Vdash \nabla \alpha \text{ iff } (\sigma(s), \alpha) \in \overline{T}(\Vdash|_{S \times \text{Base}(\alpha)}). \quad (30)$$

Example 5.12 In this example we look at the interpretation of the coalgebraic cover modality instantiated for various coalgebra types T . That is, let $\mathbb{S} = (S, \sigma, V)$ be a T -model and consider an element $\alpha \in T_\omega(\text{ML}_T)$; below we will explain what it means for the formula $\nabla\alpha$ to hold at s .

(a) In case $T = K_C$ is a constant functor, $\alpha \in T(\text{ML}_T) = C$ is just a colour $\alpha \in C$. In this case we find $\mathbb{S}, s \Vdash \nabla\alpha$ iff $\sigma(s) = \alpha$.

(b) If $T = Id$ is the identity functor, $\alpha \in T(\text{ML}_T) = \text{ML}_T$ is just a formula. We obtain the *next-time operator* of linear temporal logic: $\mathbb{S}, s \Vdash \nabla\alpha$ iff $\mathbb{S}, \sigma(s) \Vdash \alpha$.

(c) For the binary tree functor $T = Id \times Id$, the semantics of nabla is as follows: given $\alpha = (a_0, a_1) \in \text{ML}_T \times \text{ML}_T = T(\text{ML}_T)$, we have $\mathbb{S}, s \Vdash \nabla(a_0, a_1)$ iff $\mathbb{S}, t_0 \Vdash a_0$ and $\mathbb{S}, t_1 \Vdash a_1$, where t_0 and t_1 are the ‘left’ and ‘right’ successor of s , respectively, given by $\sigma(s) = (t_0, t_1)$.

(d) For the automata functor $T = 2 \times Id^C$, an element $\alpha \in T(\text{ML}_T)$ is of the form $(i, \bar{a})$ with $i \in \{0, 1\}$ and $\bar{a} = (a_c)_{c \in C}$, with each $a_c \in \text{ML}_T$. With $\sigma = \langle \chi, \tau \rangle$ we have $\mathbb{S}, s \Vdash \nabla(i, \bar{a})$ iff $\chi(s) = i$ and $\mathbb{S}, \tau(s)(c) \Vdash a_c$.

(e) Where $T = P$ is the power set functor, it is easy to verify that ∇_P is the cover modality discussed in the previous section.

Example 5.13 In this example we look at the ∇ -logic of the distribution functor D ; the nabla modality for the bag functor is defined similarly.

First of all, observe that defined as in Example B.6, D does not preserve inclusions. We can remedy this by taking a variant D' of D which takes a set S to the collection $D'S$ of *partial maps* from S to $(0, 1]$. We will not pursue this road further, but we may use it to observe that an element $\alpha \in D_\omega(\text{ML}_D)$ can be represented as a finite set $\{(a_1, p_n), \dots, (a_n, p_n)\}$ such that $p_i > 0$ for all i , and $\sum_i p_i = 1$.

For the definition of relation lifting $\overline{D}$, consider a relation $R \subseteq X_0 \times X_1$. We claim that the relation $\overline{D}R \subseteq DX_0 \times DX_1$ consists of those pairs (μ_0, μ_1) for which there is a distribution $\rho : R \rightarrow [0, 1]$ satisfying the ‘magic square conditions’:

$$\begin{aligned} &\text{for all } x_0 \in X_0. \quad \mu_0(x_0) = \sum_{y_1 \in X_1} \rho(x_0, y_1), \\ &\text{and for all } x_1 \in X_1. \quad \mu_1(x_1) = \sum_{y_0 \in X_0} \rho(y_0, x_1). \end{aligned}$$

For a concrete example, consider the table below:

	a	b	c	d	e	
1	$\square$	.17	$\square$	.08	.08	.33
2	.10	.03	.12	$\square$	.08	.33
3	$\square$	.00	.18	.12	.04	.34.
	.10	.20	.30	.20	.20	

Here we have $X_0 = \{a, b, c, d, e\}$, $X_1 = \{1, 2, 3\}$ and $R \subseteq X_0 \times X_1$ consists of those pairs of which the corresponding entry in the table is marked by a number (that is, not by $\square$). The distributions μ_0 and μ_1 are given, respectively, in the bottom row and rightmost column of the table. The entries in the table represent a distribution ρ that meets the magic square conditions.

Now let $\mathbb{S} = (S, \sigma, V)$ be a D -model, and consider a formula $\nabla\alpha$ with $\alpha = \{(a_1, p_n), \dots, (a_n, p_n)\}$. Then we find that

$\mathbb{S}, s \Vdash \nabla\alpha$ iff there is a relation $R \subseteq \Vdash$ and a map $\rho : R \rightarrow [0, 1]$ such that

$$\text{for all } i : \sum_{\{t | \mathbb{S}, t \Vdash a_i\}} \rho(t, a_i) = p_i,$$

$$\text{and for all } t \in S : \sum_{\{i | \mathbb{S}, t \Vdash a_i\}} \rho(t, a_i) = \sigma(s)(t).$$

5.3 Basic properties of ∇

In this section we prove some of the basic properties of the coalgebraic cover modality. We start with showing that ML_T is a finitary logic indeed, i.e., that every formula has only finitely many subformulas. The key property of finitary functors that will make this possible is that for every $\alpha \in T_\omega A$ there is a *smallest* subset $A' \subseteq A$ that *supports* α , i.e., is such that $\alpha \in T_\omega A'$.

Definition 5.14 Given a finitary functor T and an element $\alpha \in TX$, we define

$$\text{Base}_X^T(\alpha) := \bigcap \{Y \subseteq_\omega X \mid \alpha \in TY\}.$$

We write Base^T rather than Base^{T_ω} , and in fact omit the superscript whenever possible. $\triangleleft$

Example 5.15 The following examples are easy to check: $\text{Base}_X^{\text{Id}} : X \rightarrow P_\omega X$ is the singleton map, $\text{Base}_X^P : P_\omega X \rightarrow P_\omega X$ is the identity map on $P_\omega X$, $\text{Base}_X^{\text{Id}^2} : X \times X \rightarrow P_\omega X$ maps the pair (x_1, x_2) to the set $\{x_1, x_2\}$, and Base^D maps a finitary distribution μ to its support $\{s \in S \mid \mu(s) > 0\}$.

Proposition 5.16 Let $T : \text{Set} \rightarrow \text{Set}$ be smooth and standard. Then the following hold:

- (1) for any $\alpha \in T_\omega X$, $\text{Base}_X^T(\alpha)$ is the smallest set supporting α ;
- (2) Base^T provides a natural transformation $\text{Base} : T_\omega \rightarrow P_\omega$.

Recall that Base^T being a natural transformation means that the following diagram commutes, for any map $f : X \rightarrow Y$:

$$\begin{array}{ccc} X & T_\omega X & \xrightarrow{\text{Base}_X^T} P_\omega X \\ f \downarrow & Tf \downarrow & \downarrow P_\omega f \\ X' & T_\omega X' & \xrightarrow{\text{Base}_{X'}^T} P_\omega X' \end{array} \quad (31)$$

Proof. Part (1) of the Proposition is an immediate consequence of Proposition B.33.

For the second part, consider a map $f : X \rightarrow X'$. We have to show $P_\omega f \circ \text{Base}_X = \text{Base}_{X'} \circ T_\omega f$. Fix $\alpha \in T_\omega X$ and write $B = \text{Base}_X(\alpha)$ and $B' = \text{Base}_{X'}(T_\omega f(\alpha))$. We need

to prove $B' = f[B]$. For the inclusion $\subseteq$, first note that the following diagram must commute (recall that hooked arrows denote inclusions):

$$\begin{array}{ccc} T_\omega B & \xrightarrow{T_\omega f|_B} & T_\omega(f[B]) \\ \downarrow & & \downarrow \\ T_\omega X & \xrightarrow{T_\omega f} & T_\omega X' \end{array}$$

From this it follows that $T_\omega f(\alpha)$ actually belongs to $T_\omega(f[B])$. In other words, $f[B]$ supports $T_\omega f(\alpha)$ and, as B' is by definition the smallest such set, it follows that $B' \subseteq f[B]$. •

For the opposite inclusion $\supseteq$, we consider the diagram below. Here 1 denotes some arbitrary singleton set, and the arrow α denotes any function from 1 to $T_\omega X$ mapping the unique element of 1 to α .

$$\begin{array}{ccc} 1 & & \\ & \searrow^{T_\omega f(\alpha)} & \\ & \text{---} T_\omega(f^{-1}(B')) \text{---} & T_\omega(B') \\ & \downarrow & \downarrow \\ & T_\omega X & \xrightarrow{T_\omega f} T_\omega X' \\ & \nwarrow_{\alpha} & \end{array}$$

But since T_ω preserves weak pullbacks, the dotted arrow exists and shows that $\alpha \in T_\omega(f^{-1}(B'))$. By minimality of the base, it follows $B \subseteq f^{-1}(B')$, that is, $B' \supseteq f[B]$. QED

By Fact 5.16(1) we may find for any formula $\nabla\alpha$ a *smallest* (and finite) collection X of formulas such that $\alpha \in T_\omega X$, namely, the set $X = \text{Base}(\alpha)$. This means that we can define a natural notion of *subformula*.

Definition 5.17 We define the set $\mathbf{Sfor}(a)$ of *subformulas* of a formula $a \in \text{ML}_T$ by the following induction:

$$\begin{aligned} \mathbf{Sfor}(a) &:= \{a\} && \text{if } a \in \{p, \perp\} \\ \mathbf{Sfor}(\neg a) &:= \{\neg a\} \cup \mathbf{Sfor}(a) \\ \mathbf{Sfor}(a_0 \vee a_1) &:= \{a_0 \vee a_1\} \cup \mathbf{Sfor}(a_0) \cup \mathbf{Sfor}(a_1) \\ \mathbf{Sfor}(\nabla\alpha) &:= \{\nabla\alpha\} \cup \bigcup \{\mathbf{Sfor}(a) \mid a \in \text{Base}(\alpha)\} \end{aligned}$$

The elements of $\text{Base}(\alpha)$ will be called the *immediate* subformulas of $\nabla\alpha$. $\triangleleft$

The next properties that we consider are invariance and expressivity.

Theorem 5.18 *For any smooth and standard functor T , the language ML_T is invariant: Given any two pointed T -models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \simeq_T (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \equiv_T (\mathbb{S}', s'). \quad (32)$$

Proof. Given the smoothness of T , it suffices to prove that *bisimilarity* implies modal equivalence. Assuming that $Z : S \rightleftharpoons S'$, we will prove by induction on the complexity of ML_T -formulas that, for every ML_T -formula a :

$$\mathbb{S}, s \Vdash a \text{ iff } \mathbb{S}', s' \Vdash a, \quad (33)$$

for every pair of states $(s, s') \in Z$. Clearly this suffices to prove the proposition.

Skipping the routine parts of the proof (i.e., the base step and boolean cases of the inductive step), we focus on the case where $a = \nabla\alpha$. We only prove the direction from right to left of (33).

So, assume that $(s, s') \in Z$ and $\mathbb{S}', s' \Vdash \nabla\alpha$, and let $\Vdash \subseteq S \times \text{ML}_T$ and $\Vdash' \subseteq S' \times \text{ML}_T$ denote the satisfaction relations on $\mathbb{S}$ and $\mathbb{S}'$, respectively. It follows from $(s, s') \in Z$ that $(\sigma(s), \sigma'(s')) \in \bar{T}Z$, and from $\mathbb{S}', s' \Vdash \nabla\alpha$ that $(\sigma'(s'), \alpha) \in \bar{T}(\Vdash')$; but from the latter fact, together with the observation that $\alpha \in T\text{Base}(\alpha)$, we may derive that $(\sigma'(s'), \alpha) \in \bar{T}(\Vdash' \upharpoonright_{S \times \text{Base}(\alpha)})$ (cf. Fact B.38). Putting these observations together with the fact that $\bar{T}$ preserves relation composition, we find that

$$(\sigma(s), \alpha) \in \bar{T}(Z; \Vdash' \upharpoonright_{S \times \text{Base}(\alpha)}).$$

But by the inductive hypothesis we obtain that $Z; \Vdash' \upharpoonright_{S' \times \text{Base}(\alpha)} \subseteq \Vdash$, so that it follows by the monotonicity of relation lifting that $(\sigma(s), \alpha) \in \bar{T}(\Vdash)$. From this it is immediate by the semantics of ∇ that $\mathbb{S}, s \Vdash \nabla\alpha$, as required. QED

As could be expected, the converse of this proposition only holds if we restrict attention to *image-finite* coalgebras.

Definition 5.19 A T -coalgebra $\mathbb{S} = (S, \sigma)$ is *image-finite* if $\sigma(s) \in T_\omega S$, for all $s \in S$. $\triangleleft$

Theorem 5.20 For any smooth and standard functor T , the language ML_T is expressive on the class of image-finite T -models: Given any two pointed T_ω -models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have

$$(\mathbb{S}, s) \equiv_T (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \simeq_T (\mathbb{S}', s'). \quad (34)$$

Proof. It suffices to show that the relation of modal equivalence is itself a *bisimulation*, when restricted to the class of image-finite coalgebra models.

Fix two T -models $\mathbb{S} = (S, \sigma, V)$ and $\mathbb{S}' = (S', \sigma', V')$, and let $\equiv \subseteq S \times S'$ denote the relation of modal equivalence between $\mathbb{S}$ and $\mathbb{S}'$. (That is, we avoid notational clutter and write $\equiv$ instead of $\equiv_T$.) We will use Theorem 3.8 in order to prove that $\equiv$ is a bisimulation, and suppose for contradiction that $s \equiv s'$. Our goal will be to show that

$$(\sigma s, \sigma' s') \in \bar{T}(\equiv). \quad (35)$$

First of all, it follows by image-finiteness that we may define the (finite) sets $B := \text{Base}(\sigma s)$ and $B' := \text{Base}(\sigma' s')$. Since $\bar{T}$ commutes with restrictions it suffices to prove that

$$(\sigma s, \sigma' s') \in \bar{T}(\equiv \upharpoonright_{B \times B'}). \quad (36)$$

Furthermore, since B and B' are finite, for every $t \in B$ we may find a formula c_t such that, for all $t' \in B'$,

$$t \equiv t' \text{ iff } \mathbb{S}', t' \Vdash c_t,$$

as the reader can easily verify. But with $H := \{c_t \mid t \in B\}$, we may think of c as a surjection $c : B \rightarrow H$ satisfying $\text{Gr}(c) \subseteq \Vdash$ and

$$\equiv \upharpoonright_{B \times B'} = \text{Gr}(c) ; (\Vdash_{B' \times H})^\smile.$$

From this we may conclude, using various properties of relation lifting, that

$$\bar{T}(\equiv \upharpoonright_{B \times B'}) = \text{Gr}(Tc) ; (\bar{T} \Vdash_{B' \times H})^\smile. \quad (37)$$

Now let

$$\gamma := (Tc)(\sigma s),$$

then by definition we have $(\sigma s, \gamma) \in \text{Gr}(Tc)$. Furthermore, from $\text{Gr}(c) \subseteq \Vdash$ we derive $\text{Gr}(Tc) \subseteq \bar{T} \Vdash$, so that we find $(\sigma s, \gamma) \in \bar{T} \Vdash$. This implies that the formula $\nabla \gamma$ is *true* at s , and since $s \equiv s'$, it is also true at s' . By the semantics of the ∇ modality this means that $(\sigma s', \gamma) \in \bar{T} \Vdash$. Using various properties of relation lifting we find that $(\gamma, \sigma s') \in (\bar{T} \Vdash_{B' \times H})^\smile$. We may then conclude that $(\sigma s, \sigma s') \in \text{Gr}(Tc) ; (\bar{T} \Vdash_{B' \times H})^\smile$, which by (37) suffices to prove (36), and thus (35). In other words, we are done. QED

The last basic property that we mention is that of *satisfiability reduction*.

Proposition 5.21 *Let $\nabla \alpha$ be a formula in ML_T . Then $\nabla \alpha$ is satisfiable iff every $a \in \text{Base}(\alpha)$ is satisfiable.*

Proof. For the direction from right to left, assume that every $a \in \text{Base}(\alpha)$ is satisfiable. That is, assume that for every $a \in \text{Base}(\alpha)$ there is a pointed model $(\mathbb{S}_a, s_a)$, with $\mathbb{S}_a = (S_a, \sigma_a, V_a)$ and such that $\mathbb{S}_a, s_a \Vdash a$, for each a . We define a new model $\mathbb{S} = (S, \sigma, V)$, where $S := \{r\} \uplus \biguplus \{S_a \mid a \in \text{Base}(\alpha)\}$. For the valuation V we simply define $V(p) := \bigcup_a V_a(p)$, while on an element $s \in S_a$ the coalgebra map $\sigma : S \rightarrow TS$ is defined by putting $\sigma(s) := \sigma_a(s)$. For the definition of the unfolding $\sigma(r)$ of the ‘root’ r , consider the map $f : \text{Base}(\alpha) \rightarrow S$ given by $a \mapsto s_a$, and simply put $\sigma(r) := (Tf)(\alpha)$. It is then immediate by the definition of f that $(\text{Gr}f)^\smile \subseteq \Vdash$, so that we find, using various properties of relation lifting (cf. Fact B.36 and B.38):

$$(\sigma r, \alpha) \in (\text{Gr}(Tf))^\smile = \bar{T}((\text{Gr}f)^\smile) \subseteq \bar{T}(\Vdash),$$

from which it follows that $\mathbb{S}, r \Vdash \nabla \alpha$ indeed.

For the opposite direction, we need the following little fact about the functor T :

$$\text{for any } f : A \rightarrow B \text{ and any } \alpha \in TA \text{ we have } (Tf)(\alpha) \in T(f[A]). \quad (38)$$

To see why (38) holds, factorize f as the unique composition $f = \iota \circ f'$ of a surjection $f' : A \twoheadrightarrow f[A]$ and an inclusion $\iota : f[A] \hookrightarrow B$. From this factorization it follows that $Tf = (T\iota) \circ (Tf')$, where $f' : TA \twoheadrightarrow Tf[A]$ is surjective since every set functor preserves surjections,

and $T\iota : Tf[A] \hookrightarrow TB$ is an inclusion by assumption on T . From these observations (38) is easy to derive.

Now assume that $\nabla\alpha$ is satisfiable, then there is a T -model $\mathbb{S} = (S, \sigma, V)$ such that $\mathbb{S}, s \Vdash \nabla\alpha$ for some state $s \in S$. Then by definition of the semantics of ∇ we have that $(\sigma s, \alpha) \in \bar{T}(\Vdash)$, and so by definition of $\bar{T}$ there is an object $\rho \in T(\Vdash)$ such that $T\pi_0(\rho) = \sigma(s)$ and $T\pi_1(\rho) = \alpha$, where $\pi_0 : \Vdash \rightarrow S$ and $\pi_1 : \Vdash \rightarrow \mathbf{ML}_T$ are the projection functions on the relation $\Vdash$. But then it follows by (38) that $\alpha \in T(\pi_1[\Vdash]) = T(\text{Ran}(\Vdash))$, so that $\text{Base}(\alpha) \subseteq \text{Ran}(\Vdash)$. In other words, for every $a \in \text{Base}(\alpha)$ there is an $s \in S$ where a holds. In particular, this means that every $a \in \text{Base}(\alpha)$ is satisfiable. QED

5.4 Coalgebraic modal distributive laws

In this section we will formulate three *coalgebraic modal distributive laws (CMDLs)* describing the interaction between the coalgebraic modality ∇ on the one hand, and the boolean operations on the other. For a concise formulation of these principles it will be convenient to slightly rearrange the coalgebraic modal language, working with the finitary versions $\bigwedge$ and $\bigvee$ of the binary connectives for conjunction and disjunction. That is, in this section we will be working with the following variant of the language.

Definition 5.22 The language L_T is given by the following grammar:

$$a ::= p \mid \neg a \mid \bigwedge A \mid \bigvee A \mid \nabla\alpha$$

where $p \in \mathbf{Q}$, $A \in P_\omega L_T$ and $\alpha \in T_\omega L_T$. $\triangleleft$

The semantics of $\bigwedge$ and $\bigvee$ is as expected:

$$\begin{aligned} \mathbb{S}, s \Vdash \bigwedge A & \quad \text{iff} \quad \mathbb{S}, s \Vdash a \text{ for all } a \in A \\ \mathbb{S}, s \Vdash \bigvee A & \quad \text{iff} \quad \mathbb{S}, s \Vdash a \text{ for some } a \in A \end{aligned}$$

In particular, this means that instead of taking $\top$ and $\perp$ to be primitive symbols, we may consider them as abbreviations:

$$\begin{aligned} \top &:= \bigwedge \emptyset \\ \perp &:= \bigvee \emptyset. \end{aligned}$$

A key aspect of the formulation of the CMDLs is the observation that we may think of the connectives $\bigwedge, \bigvee$ and $\neg$ as *maps* of the respective types $\bigwedge, \bigvee : P_\omega L_T \rightarrow L_T$ and $\neg : L_T \rightarrow L_T$. In particular, this perspective allows us to apply the functor T to these connectives, obtaining maps $T\bigwedge, T\bigvee : T_\omega P_\omega L_T \rightarrow T_\omega L_T$, and $T\neg : T_\omega L_T \rightarrow T_\omega L_T$. Thus, for any object $\Phi \in T_\omega P_\omega L_T$ we find $(T\bigvee)\Phi \in T_\omega L_T$, which means that $\nabla(T\bigvee)\Phi$ is a well-formed formula.

Convention 5.23 Since we will be dealing here with formulas and similar objects in various, closely related sets, including $\mathbf{Q}, L_T, T_\omega L_T, P_\omega L_T, P_\omega T_\omega L_T$ and $T_\omega P_\omega L_T$, it will be convenient to use some kind of *naming convention*, see Table 1.

In order to formulate the modal distributive laws we need some auxiliary definitions.

Set	Elements
$\mathbf{Q}$	$p, q, \dots$
L_T	$a, b, \dots$
$T_\omega L_T$	$\alpha, \beta, \dots$
$P_\omega L_T$	$A, B, \dots$
$P_\omega T_\omega L_T$	$\Gamma, \Delta, \dots$
$T_\omega P_\omega L_T$	$\Phi, \Psi, \dots$

Table 1: Naming convention

Definition 5.24 Given a smooth and standard set functor T , we define, for every set X , a function $\lambda_X^T : TPX \rightarrow PTX$ by putting

$$\lambda_X^T(\Phi) := \{\alpha \in TX \mid (\alpha, \Phi) \in \bar{T}(\in_X)\}, \quad (39)$$

where $\in_X$ denotes the membership relation $\in$, restricted to $X \times PX$. Elements of $\lambda_X^T(\Phi)$ will be referred to as *lifted members* of Φ . The family $\lambda^T = \{\lambda_X^T\}_{X \in \mathbf{Set}}$ will be called the *T-transformation*.

A set $\Phi \in TPX$ is a *redistribution* of a set $\Gamma \in PTX$ if $\Gamma \subseteq \lambda_X^T(\Phi)$, that is, every element of Γ is a lifted member of Φ . In case $\Gamma \in P_\omega T_\omega X$, we call a redistribution Φ *slim* if $\Phi \in T_\omega P_\omega(\bigcup_{\gamma \in \Gamma} \text{Base}(\gamma))$. The set of slim redistributions of Γ is denoted as $SRD(\Gamma)$. $\triangleleft$

Remark 5.25 Properties of $\bar{T}$ are intimately related to those of λ^T . In order to express the connection, we need to introduce the concept of a (categorical) distributive law.

Let T be a covariant set functor. A *distributive law* of T over a (co- or contravariant) set functor M is a natural transformation $\theta : TM \rightarrow MT$; that is, the following diagram commutes, for every map $f : X \rightarrow Y$:

$$\begin{array}{ccccc} X & & TMX & \xrightarrow{\theta_X} & MTX \\ f \downarrow & & TMf \downarrow & & \downarrow MTf \\ Y & & TMY & \xrightarrow{\theta_Y} & MTY \end{array}$$

(In case M is a covariant functor the downward arrows have to be reversed.)

Now observe that we can in fact *define* the family of functions $\lambda^T = \{\lambda_X^T\}_{X \in \mathbf{Set}}$ as in (39) for any set function T . If T preserves weak pullbacks, $\lambda^T = \{\lambda_X^T\}_{X \in \mathbf{Set}}$ is a distributive law of T over the power set functor P .

Definition 5.26 Let T be a smooth and standard set functor which restricts to finite sets. Consider the following *coalgebraic modal distributive laws*:

$$\begin{aligned} (DL_\vee) \quad \nabla(T \vee)(\Phi) &\equiv \bigvee \{ \nabla \alpha \mid (\alpha, \Phi) \in \bar{T}(\in_{L_T}) \} \\ (DL_\wedge) \quad \bigwedge \{ \nabla \gamma \mid \gamma \in \Gamma \} &\equiv \bigvee \{ \nabla(T \wedge)(\Phi) \mid \Phi \in SRD(\Gamma) \} \\ (DL_\neg) \quad \neg \nabla \alpha &\equiv \bigvee \{ \nabla T(\wedge \circ P \neg) \Psi \mid \Psi \in T_\omega P_\omega \text{Base}(\alpha) \text{ and } (\alpha, \Psi) \notin \bar{T}(\notin) \} \end{aligned} \quad \triangleleft$$

Note that the restriction to finite sets is needed in order to keep the disjunctions on the right hand side of the equivalences finite

Proposition 5.27 *Let T be a smooth and standard set functor which restricts to finite sets. All three coalgebraic modal distributive laws are valid.*

Proof. In order to understand the validity of these laws, fix some T -model $\mathbb{S} = (S, \sigma, V)$.

We first consider (DL_V) , proving the direction from left to right. First observe that for any $A \subseteq_\omega L_T$ we have $\mathbb{S}, s \Vdash \bigvee A$ iff $\mathbb{S}, s \Vdash a$, for some $a \in A$. Putting it differently, the relations $\Vdash; \in$ and $\Vdash; \bigvee^\circ$ coincide⁹. From this it follows that

$$\overline{T}(\Vdash; \in) = \overline{T}(\Vdash; \bigvee^\circ). \quad (40)$$

Now fix some object $\Phi \in T_\omega P_\omega L$, and suppose that s is a state in $\mathbb{S}$ such that $s \Vdash \nabla(T \bigvee) \Phi$. By the truth definition, the pair $(\sigma(s), (T \bigvee)(\Phi))$ belongs to the relation $\overline{T}(\Vdash)$, and so $(\sigma(s), \Phi)$ belongs to $(\overline{T} \Vdash); (T \bigvee)^\circ = \overline{T}(\Vdash; \bigvee^\circ)$. But then by (40), we find $(\sigma(s), \Phi) \in \overline{T}(\Vdash; \in) = \overline{T} \Vdash; \overline{T} \in$. In other words, there is some object β such that $(\sigma(s), \beta) \in \overline{T}(\Vdash)$ and $(\beta, \Phi) \in \overline{T}(\in)$. Clearly then $s \Vdash \nabla \beta$, and so we have $s \Vdash \bigvee \{\nabla \beta \mid \beta \overline{T} \in \Phi\}$, as required.

For the validity of $(DL_\wedge)$, we also confine attention to the direction from left to right. Assume that $\mathbb{S}, s \Vdash \nabla \gamma$ for all $\gamma \in \Gamma$. We need to come up with some slim redistribution Φ of Γ such that $\mathbb{S}, s \Vdash \nabla(T \wedge) \Phi$. For this purpose we associate, with any state $t \in S$, the finite set

$$A_t := \{a \in \bigcup_{\gamma \in \Gamma} \text{Base}(\gamma) \mid \mathbb{S}, t \Vdash a\}.$$

Taking A to be the map $A : S \rightarrow P_\omega L_T$, we may define $\Phi := (TA)(\sigma(s)) \in T_\omega P_\omega L_T$.

First we show that $\mathbb{S}, s \Vdash \nabla(T \wedge) \Phi$. Observe that by definition of the map $A : S \rightarrow P_\omega L_T$, the function $\bigwedge \circ A : S \rightarrow L_T$ is such that

$$\text{Gr}(\bigwedge \circ A) \subseteq \Vdash.$$

From this we obtain

$$\text{Gr}((T \bigwedge) \circ (TA)) \subseteq \overline{T}(\Vdash)$$

by the properties of the operation $\overline{T}$. But that means that for every element $\tau \in TS$, we have that $(\tau, ((T \bigwedge) \circ (TA))(\tau)) \in \overline{T} \Vdash$. In particular, we find that $(\sigma s, (T \bigwedge) \Phi) = (\sigma s, (T \bigwedge)(TA)(\sigma(s))) \in \overline{T} \Vdash$, showing that $\mathbb{S}, s \Vdash \nabla(T \bigwedge) \Phi$ as required.

It is left to prove that Φ is a slim redistribution of Γ . Observe that by definition of the map A , we have that

$$\text{Gr}(A); \in^\circ = \Vdash|_{S \times B},$$

where $B := \bigcup_{\gamma \in \Gamma} \text{Base}(\gamma)$. From this it follows by the properties of relation lifting that

$$\text{Gr}(TA); (\overline{T} \in)^\circ = \overline{T}(\Vdash)|_{TS \times TB}.$$

⁹Here we write $\bigvee$ instead of $\text{Gr}(\bigvee)$

But then for each $\gamma \in \Gamma$ we may derive from the fact that $(\sigma s, \gamma) \in \overline{T}(\Vdash) \upharpoonright_{TS \times TB}$ that there is some object Ψ such that $(\sigma s, \Psi) \in \text{Gr}(TA)$ and $(\Psi, \gamma) \in (\overline{T}\in)^\circ$. It then easily follows that $\Psi = (TA)(\sigma s) = \Phi$ and so $(\gamma, \Phi) = (\gamma, \Psi) \in \overline{T}(\in)$. In other words, each $\gamma \in \Gamma$ is a lifted member of Φ , and so Φ is a redistribution of Γ ; but then by its definition it is slim.

Finally, the validity of $(DL_{\neg})$ is left as an exercise to the reader.

QED

5.5 Coalgebraic Logic

We will now see that the coalgebraic modal distributive laws that we proved in the previous section are in fact quite strong principles, with important applications.

We start with the coalgebraic generalisation of the disjunctive normal form result on the cover modality, Theorem 5.8.

Definition 5.28 We let the following grammar:

$$a ::= p \mid \neg p \mid \perp \mid \top \mid a_0 \vee a_1 \mid \pi \bullet \nabla \alpha.$$

define the set $\text{DML}_T(\mathbf{Q})$ of *disjunctive T -modal formulas* in $\mathbf{Q}$. Here we have $p \in \mathbf{Q}$ and π is a finite conjunction of literals over $\mathbf{Q}$. $\triangleleft$

The proof of the following theorem is completely analogous to that of Theorem 5.8.

Theorem 5.29 *Let T be a smooth and standard set functor which restricts to finite sets. The languages ML_T and DML_T are effectively equi-expressive.*

For the following result recall that a modal logic $(L, \Vdash)$ has the *finite model property* if every satisfiable L -formula is satisfiable in a finite model.

Theorem 5.30 *Let T be a smooth and standard set functor which restricts to finite sets. Then ML_T has the finite model property.*

Proof. By Theorem 5.29 it suffices to prove the finite model property for *disjunctive* formulas. We leave it as an exercise for the reader to establish this result — this goes by a straightforward proof by induction on the complexity of DML_T -formulas, of which the inductive case for the ∇ modality uses the observation underlying the proof of our satisfiability reduction result, Proposition 5.21. QED

Remark 5.31 Theorem 5.29 can also be used to obtain *decidability* results for logics ML_T . For instance, it can be proved that the satisfiability problem for the language $\text{ML}_P = \text{ML}_{\nabla}$ can be solved in *linear time*. However, since these results depend on the functor, or more specifically: on the representation of formulas of the form $\nabla \alpha$, we refrain from going into detail here.

Remark 5.32 As another corollary of Theorem 5.29 we can show that for any smooth and standard set functor T which restricts to finite sets, the logic ML has *uniform interpolation*, a strong version of Craig's interpolation property. •

Finally, we briefly mention a sound and complete derivation system for the set of valid ML-formulas.

Definition 5.33 Let T be a smooth and standard set functor which restricts to finite sets. For the *derivation system* $\mathbf{M}$, we start with fixing an arbitrary sound and complete set of axioms and rules for classical propositional logic; we extend this with the following derivation rule:

$$\frac{\{a \rightarrow b \mid (a, b) \in Z\}}{\nabla\alpha \rightarrow \nabla\beta} \quad (\alpha, \beta) \in \overline{T}Z,$$

together with the one-sided versions of the coalgebraic modal distributive laws:

$$\begin{aligned} (A_{\vee}) \quad & \nabla(T \vee)(\Phi) \rightarrow \bigvee \{ \nabla\alpha \mid (\alpha, \Phi) \in \overline{T}(\in_X) \} \\ (A_{\wedge}) \quad & \bigwedge \{ \nabla\gamma \mid \gamma \in \Gamma \} \rightarrow \bigvee \{ \nabla(T \wedge)(\Phi) \mid \Phi \in SRD(\Gamma) \} \\ (A_{\neg}) \quad & \neg\nabla\alpha \rightarrow \bigvee \{ \nabla T(\wedge \circ P\neg)\Psi \mid \Psi \in T_{\omega}P_{\omega}Base(\alpha) \text{ and } (\alpha, \Psi) \notin \overline{T}(\notin) \} \end{aligned} \quad \triangleleft$$

► Formulate completeness

6 Coalgebraic modalities via predicate liftings

In this chapter we take an approach to coalgebraic modal logic where the modalities are in 1-1 correspondence with so-called *predicate liftings* for the functor T . That is, with each set Λ of such predicate liftings we will associate a modal formalism ML_Λ for T -coalgebras. As a result this set-up is not completely uniform in the coalgebra type T , but it has some advantages over the approach based on relation lifting. First of all, the language of ML_Λ is completely standard, with a syntax that adds to the language of propositional logic an n -ary modality $\heartsuit_\lambda$ for each n -ary predicate lifting $\lambda \in \Lambda$. Second, there is no reason to restrict attention to functors that are smooth (preserve weak pullbacks). And finally, predicate liftings provide a uniform framework to many well-known variants of standard modal logic (including monotone and probabilistic modal logic, which were already mentioned in section 1.5).

Before we introduce the approach in full generality, we briefly discuss a few other concrete variants of standard modal logic that are covered by the approach.

6.1 Variants of modal logic

Example 6.1 (1) The *next-time operator* $\bigcirc$ of *linear time logic* is perhaps the most simple example. For its definition, consider models of the form (ω, V) , where $V : \mathbb{Q} \rightarrow P(\omega)$ is a valuation on the set ω of natural numbers; the modality $\bigcirc$ is interpreted as follows:

$$\omega, V, n \Vdash \bigcirc \varphi \text{ iff } \omega, V, n+1 \Vdash \varphi.$$

Clearly the semantics of this operator can be generalised to arbitrary T -models for the identity functor $T = \text{Id}$.

(2) Similarly, on the binary tree 2^ω we can interpret two modalities $\bigcirc_0$ and $\bigcirc_1$, with the following interpretation:

$$2^\omega, V, u \Vdash \bigcirc_i \varphi \text{ iff } 2^\omega, V, u \cdot i \Vdash \varphi,$$

where $i \in \{0, 1\}$ and $V : \mathbb{Q} \rightarrow P(2^\omega)$ is a valuation on the set of finite words over the alphabet $2 = \{0, 1\}$.

The semantics of these operators can be generalised to arbitrary models for the binary tree functor $T = \text{Id} \times \text{Id}$.

(3) Graded modal logic is a version of modal logic that allows statements about the *number* of successors that satisfy a certain formula. Formally, interpreted in Kripke models, the modality $\Diamond_{\geq k}$ has the following semantics:

$$\mathbb{S}, n \Vdash \Diamond_{\geq k} \varphi \text{ iff } s \text{ has at least } k \text{ } \varphi\text{-successors,}$$

where a φ -successor of s is a state $t \in R(s)$ where φ holds. If we restrict attention to image-finite Kripke models, it also makes sense to introduce the following ‘majority modality’ M :

$$\mathbb{S}, n \Vdash M(\varphi, \psi) \text{ iff } s \text{ has more } \varphi\text{-successors than } \psi\text{-successors.}$$

Note that these modalities are not bisimulation invariant if we consider Kripke frames as coalgebras for the powerset functor. However, as we will see below, we may also see Kripke

frames as coalgebras for the *bag* functor B (see the appendix for its definition), and for that functor both modalities will turn out to be invariant.

As we will see in this section, the common semantic pattern in many of these formalisms can be captured rather nicely in a coalgebraic framework by the notion of a *predicate lifting*.

6.2 Modalities via predicate liftings

To introduce the notion of a predicate lifting, we consider the example of probabilistic modal logic. In Example 1.8 we defined the semantics of the modality $\Diamond_q$ (with q a rational number in $[0, 1]$) in a model $\mathbb{S} = (S, \sigma, V)$ for the distribution functor D , as follows:

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sum_{u \in \llbracket \varphi \rrbracket^{\mathbb{S}}} \sigma(s)(u) > q, \quad (41)$$

where we recall that $\llbracket \varphi \rrbracket^{\mathbb{S}}$ denotes the *extension* of φ in $\mathbb{S}$, that is, the set $\llbracket \varphi \rrbracket = \{t \in S \mid \mathbb{S}, t \Vdash \varphi\}$ of states in $\mathbb{S}$ where φ is true. The way that we will be thinking of this definition now is as

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sigma(s) \in \left\{ \mu \in D(S) \mid \sum_{u \in \llbracket \varphi \rrbracket^{\mathbb{S}}} \mu(u) > q \right\}, \quad (42)$$

or, in fact, as

$$\mathbb{S}, s \Vdash \Diamond_q \varphi \text{ iff } \sigma(s) \in \theta_S^q(\llbracket \varphi \rrbracket^{\mathbb{S}}), \quad (43)$$

where $\theta_S^q : PS \rightarrow PDS$ is defined by

$$\theta_S^q : U \mapsto \left\{ \mu \in D(S) \mid \sum_{u \in U} \mu(u) > q \right\}.$$

In other words, we may think of the semantics of the modality $\Diamond_q$ as being indexed by a family θ^q of maps $\theta_S^q : PS \rightarrow PDS$, where each θ_S^q *lifts* a predicate on S (i.e., a subset of S) to a predicate on DS .

Now in principle we may associate a modality with each such family θ . However, as we will see below, it will make a lot of sense to impose the following uniformity condition on the family of maps: We will require that, for each map $f : S' \rightarrow S$, the following diagram commutes:

$$\begin{array}{ccccc} S' & & PS' & \xrightarrow{\theta_{S'}} & PDS' \\ f \downarrow & & \uparrow \check{P}f & & \uparrow \check{P}Df \\ S & & PS & \xrightarrow{\theta_S} & PDS \end{array}$$

That is, we will require a ‘proper’ predicate lifting for the distribution functor to be a *natural transformation* $\theta : \check{P} \rightarrow \check{P}D$, where $\check{P}$ is the *contravariant* powerset functor. In general, for an arbitrary set functor T we introduce the concept of a *predicate lifting* of some arbitrary but fixed finite arity, as follows.

Definition 6.2 A *predicate lifting* is a natural transformation of the form $\lambda : \check{P}^n \rightarrow \check{P}T$, for some number $n \in \omega$ which we shall refer to as the *arity* of λ , notation: $n = \text{ar}\lambda$. $\triangleleft$

Recall that the *naturality condition* on predicate liftings means that the following diagram commutes, for every function $f : S' \rightarrow S$:

$$\begin{array}{ccc} S' & (PS')^n \xrightarrow{\lambda_{S'}} & PTS' \\ f \downarrow & \uparrow (\check{P}f)^n & \uparrow \check{P}Tf \\ S & (PS)^n \xrightarrow{\lambda_S} & PTS \end{array} \quad (44)$$

The idea now is that with each predicate lifting λ we may associate a modality $\heartsuit_\lambda$, of the same arity as λ , and that the semantics of $\heartsuit_\lambda$ in a coalgebra (S, σ) is defined in terms of the predicate lifting λ itself (and the coalgebra map σ). The logics that we about to introduce are thus parametrised by a collection of predicate liftings; such collections we shall refer to as *modal signatures*.

Definition 6.3 A *modal signature* for a set functor T is nothing but a set of predicate liftings for T . Given such a collection Λ , and a set $\mathbf{Q}$ of proposition letters, the formulas of the modal logic $\text{ML}_\Lambda(\mathbf{Q})$ are given by the following grammar:

$$\varphi ::= p \mid \perp \mid \neg\varphi \mid \varphi_0 \vee \varphi_1 \mid \heartsuit_\lambda(\varphi_0, \dots, \varphi_{n-1}),$$

where $p \in \mathbf{Q}$, and λ is an n -ary predicate lifting for T . $\triangleleft$

We will also use the connectives $\wedge$, $\rightarrow$ and $\leftrightarrow$ as the standard abbreviations.

The semantics of the languages ML_Λ is defined in a uniform way, with the modality $\heartsuit_\lambda$ being interpreted ‘by λ itself’.

Definition 6.4 Let $\mathbb{S} = (S, \sigma, V)$ be a T -model for some set functor T . We define the satisfaction relation $\Vdash_{\mathbb{S}} \subseteq S \times \text{ML}_\Lambda(\mathbf{Q})$ by induction on the complexity of $\text{ML}_\Lambda(\mathbf{Q})$ -formulas. With all other clauses of this definition being standard, we only mention the clause for the coalgebraic modalities:

$$\mathbb{S}, s \Vdash \heartsuit_\lambda(\varphi_0, \dots, \varphi_{n-1}) \text{ iff } \sigma(s) \in \lambda(\llbracket \varphi_0 \rrbracket^{\mathbb{S}}, \dots, \llbracket \varphi_{n-1} \rrbracket^{\mathbb{S}}). \quad (45)$$

The notions of satisfiability, validity, and ML_Λ -equivalence are all defined in the obvious way (cf. Definition 5.10 for the ∇ -logic); the relation of ML_Λ -equivalence between formulas will usually be denoted as $\equiv_\Lambda$ rather than as $\equiv_{\text{ML}_\Lambda}$. $\triangleleft$

Remark 6.5 A succinct way of defining the semantics of the modality $\heartsuit_\lambda$ (45) is as follows:

$$\llbracket \heartsuit_\lambda(\varphi_0, \dots, \varphi_{n-1}) \rrbracket^{\mathbb{S}} := (\check{P}\sigma)(\lambda(\llbracket \varphi_0 \rrbracket^{\mathbb{S}}, \dots, \llbracket \varphi_{n-1} \rrbracket^{\mathbb{S}})). \quad (46)$$

Example 6.6 (1) The box and diamond modalities of standard modal logic can be seen as the coalgebraic modalities associated with the unary predicate liftings $\lambda^\square, \lambda^\diamond : \check{P} \rightarrow \check{P}P$ given by

$$\begin{aligned}\lambda_S^\square : U &\mapsto \{X \in PS \mid X \subseteq U\}, \\ \lambda_S^\diamond : U &\mapsto \{X \in PS \mid X \cap U \neq \emptyset\}.\end{aligned}$$

We quickly verify that $\lambda^\square$ satisfies the naturality condition (44). In order to show that $\lambda_{S'}^\square \circ \check{P}f = (\check{P}Pf) \circ \lambda_S^\square$, it suffices to show that the following identities hold, for every $U \in PS$:

$$\begin{aligned}\lambda_{S'}^\square(\check{P}f(U)) &= \{X' \in PS' \mid X' \subseteq \check{P}f(U)\} && \text{(definition } \lambda_{S'}^\square) \\ &= \{X' \in PS' \mid fx' \in U, \text{ all } x' \in X'\} && \text{(obvious)} \\ &= \{X' \in PS' \mid Pf(X') \subseteq U\} && \text{(obvious)} \\ &= \{X' \in PS' \mid Pf(X') \in \lambda_S^\square(U)\} && \text{(definition } \lambda_S^\square) \\ &= (\check{P}Pf)(\lambda_S^\square(U)). && \text{(definition } \check{P}Pf)\end{aligned}$$

(2) In the case of monotone modal logic, the box and diamond modalities are induced by the following predicate liftings $\mu^\square, \mu^\diamond : \check{P} \rightarrow \check{P}M$:

$$\begin{aligned}\mu_S^\square : U &\mapsto \{\sigma \in MS \mid U \in \sigma\}, \\ \mu_S^\diamond : U &\mapsto \{\sigma \in MS \mid (S \setminus U) \notin \sigma\}.\end{aligned}$$

(3) The next-time operator $\bigcirc$ of linear temporal logic is obtained from the *identity lifting* $\lambda^\bigcirc : \check{P} \rightarrow \check{P}$:

$$\lambda_S^\bigcirc : U \mapsto U.$$

Example 6.7 Similarly to the case of probabilistic modal logic, for each $k \in \mathbb{N}$ we can define a predicate lifting $\lambda^{\geq k} : \check{P} \rightarrow \check{P}B$ for the bag functor B :

$$\lambda_S^{\geq k} : U \mapsto \{\mu : S \rightarrow \mathbb{N}^\infty \mid \sum_{u \in U} \mu(u) \geq k\}.$$

Now consider a Kripke model $\mathbb{S} = (S, \sigma, V)$, with $\sigma : S \rightarrow PS$. We may think of σ as a coalgebra map $\sigma^\circ : S \rightarrow BS$ for the functor B , by putting

$$\sigma^\circ(s)(t) := \begin{cases} 1 & \text{if } t \in \sigma(s), \\ 0 & \text{if } t \notin \sigma(s). \end{cases}$$

It is then straightforward to verify that for any Kripke model $\mathbb{S}$ we have

$$\mathbb{S}, n \Vdash \heartsuit_{\lambda^{\geq k}} \varphi \text{ iff } s \text{ has at least } k \text{ } \varphi\text{-successors,}$$

so that we can indeed think of graded modal logic as a coalgebraic logic.

Similarly, the ‘majority modality’ M of Example 6.1(3) can be seen as the coalgebraic modality that is induced by the *binary* predicate lifting $\lambda^M : \check{P}^2 \rightarrow \check{P}B_\omega$ given by

$$\lambda_S^M : (U_0, U_1) \mapsto \{\mu : S \rightarrow \mathbb{N} \mid \sum_{u \in U_0} \mu(u) > \sum_{u \in U_1} \mu(u)\}.$$

Remark 6.8 Nullary predicate liftings exist. To unravel their meaning, note that we may think of any set of the form $(PS)^0$ as a *singleton* (more precisely, as the singleton consisting of the unique map $!_S : 0 \rightarrow PS$, where $0 = \emptyset$ is the empty set). Hence, we may identify a map $\lambda_S : (\check{P}S)^0 \rightarrow \check{P}TS$ with a *distinguished element* $\lambda_S(!_S)$ of the set $\check{P}TS$, i.e., a subset of TS , and the naturality condition states that

$$(PTf)(\lambda_{S'}(!_S)) = \lambda_S(!_S),$$

for any map $f : S \rightarrow S'$.

Now suppose that λ is such a nullary predicate lifting, then the nullary modality $\heartsuit_\lambda$ associated with λ can be seen as a modal *constant*:

$$\mathbb{S}, \sigma, s \Vdash \heartsuit_\lambda \text{ iff } \sigma(s) \in \lambda_S(!_S).$$

Below we give two natural examples of this phenomenon.

Example 6.9 A natural example of a nullary modality is the constant $\checkmark$ that is sometimes used to indicate that a state in a finite deterministic automaton is accepting. Recall that these devices are coalgebras for the functor $2 \times Id^C$, and consider the nullary predicate lifting $\lambda^\checkmark : \check{P}^0 \rightarrow \check{P}(2 \times Id^C)$ given by

$$\lambda_S^\checkmark(!_S) := \{(i, f) \in 2 \times S^C \mid i = 1\}.$$

We obtain, for any state s in a deterministic automaton $\mathbb{S} = (S, \sigma)$, that s is accepting iff $\sigma(s) \in \lambda_S^\checkmark(!_S)$, so that we may think of the predicate lifting $\lambda^\checkmark$ as inducing the modality $\checkmark$.

Example 6.10 Let T be a functor, and $\mathbf{Q}$ a set of proposition letters. Recall from Remark 1.20 that we may see a T -model (S, σ, V) over $\mathbf{Q}$ as a coalgebra (S, σ_V) for the functor $T_{\mathbf{Q}} = K_{P\mathbf{Q}} \times T$, where $\sigma_V : S \rightarrow P\mathbf{Q} \times TS$ is defined by putting

$$\sigma_V(s) := (V^b(s), \sigma(s)).$$

Now fix a proposition letter $q \in \mathbf{Q}$, and consider the following *nullary* predicate lifting $\lambda^q : \check{P}^0 \rightarrow \check{P}T_{\mathbf{Q}}$ for this functor:

$$\lambda_S^q(!_S) := \{(c, \tau) \in P\mathbf{Q} \times TS \mid q \in c\}.$$

Furthermore, observe that the *modality* associated with this predicate lifting is also nullary, that is, a constant; its semantics in a $T_{\mathbf{Q}}$ -coalgebra (S, ξ) is given by

$$\begin{aligned} \mathbb{S}, \xi, x \Vdash \heartsuit_{\lambda^q} & \quad \text{iff} \quad \xi s \in \lambda_X^q(!_X). \\ & \quad \text{iff} \quad q \in \pi_0(\xi s). \end{aligned}$$

In particular, this means that if (S, ξ) is of the form (S, σ_V) for some T -model (S, σ, V) , we obtain that

$$(S, \sigma, V), s \Vdash q \text{ iff } (S, \sigma_V), s \Vdash \heartsuit_{\lambda^q} \tag{47}$$

Based on this equivalence, we may think of proposition letters as modalities associated with nullary predicate liftings.

Remark 6.11 Given a set functor T and a set $\mathbf{Q}$ of proposition letters, we can now make the connection explicit between modal languages for T -models over $\mathbf{Q}$ on the one hand, and for $T_{\mathbf{Q}}$ -coalgebras on the other.

Based on Example 6.10, we see that there is a 1-1 connection between proposition letters in $\mathbf{Q}$ and nullary predicate liftings for $T_{\mathbf{Q}}$ that dissect the ‘PQ-part’ c of an arbitrary object $(c, \tau) \in T_{\mathbf{Q}}S$.

To finish the picture, we now associate with an arbitrary n -ary predicate lifting $\lambda : P^n \rightarrow PT$ for T , an n -ary predicate lifting $\lambda' : P^n \rightarrow PT_{\mathbf{Q}}$ for $T_{\mathbf{Q}}$ as follows:

$$\lambda'_S(X_0, \dots, X_{n-1}) := \{(c, \tau) \in T_{\mathbf{Q}} \mid \tau \in \lambda_S(X_0, \dots, X_{n-1})\}.$$

Then, given a modal signature Λ for T and a set $\mathbf{Q}$ of proposition letters, we define the signature $\Lambda + \mathbf{Q}$ for the functor $T_{\mathbf{Q}}$ by putting

$$\Lambda + \mathbf{Q} := \{\lambda' \mid \lambda \in \Lambda\} \cup \{\lambda^q \mid q \in \mathbf{Q}\},$$

and we leave it for the reader to verify that with this definition, we can see the language $\text{ML}_{\Lambda}(\mathbf{Q})$ (for T -models over $\mathbf{Q}$) and $\text{ML}_{\Lambda+\mathbf{Q}}(\emptyset)$ (for $T_{\mathbf{Q}}$ -coalgebras) as notational variants of one another. In the sequel we will use this observation and use the language $\text{ML}_{\Lambda}(\mathbf{Q})$ for $T_{\mathbf{Q}}$ -coalgebras; in particular, we will always simply write q instead of $\heartsuit_{\lambda^q}$.

6.3 Basic properties of ML_{Λ}

In this subsection we make some first observations about the modal logic of predicate liftings. First we show how the naturality condition (44) implies invariance under behavioural equivalence.

Theorem 6.12 *Let Λ be a modal signature for the set functor T . Then the language ML_{Λ} is invariant: Given any two pointed T -models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have*

$$(\mathbb{S}, s) \simeq_T (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \equiv_{\Lambda} (\mathbb{S}', s'). \quad (48)$$

Proof. Given the definition of behavioural equivalence, it suffices to prove that for any coalgebra morphism $f : \mathbb{S} \rightarrow \mathbb{S}'$, and any state $s \in S$, we have $(\mathbb{S}, s) \equiv_{\Lambda} (\mathbb{S}', fs)$. So fix $f : \mathbb{S} \rightarrow \mathbb{S}'$; we will show that every formula $\varphi \in \text{ML}_{\Lambda}(\mathbf{Q})$ satisfies the following:

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \mathbb{S}', fs \Vdash \varphi, \text{ for all states } s \in S,$$

or equivalently,

$$\llbracket \varphi \rrbracket^{\mathbb{S}} = (\check{P}f) \llbracket \varphi \rrbracket^{\mathbb{S}'}. \quad (49)$$

We will prove (49) by a straightforward formula induction. Leaving the routine cases as an exercise to the reader, we focus on the case where $\varphi = \heartsuit_{\lambda}(\psi_0, \dots, \psi_{n-1})$. Our proof makes use of the fact that the diagram below commutes. This should be obvious for the left rectangle,

witnessing the naturality of λ ; observe that the right rectangle is obtained by applying the (contravariant!) functor $\check{P}$ to the diagram indicating that $f : S \rightarrow S'$ is a coalgebra morphism.

$$\begin{array}{ccccc}
 (PS)^n & \xrightarrow{\lambda_S} & PTS & \xrightarrow{\check{P}\sigma} & PS \\
 (\check{P}f)^n \uparrow & & \uparrow \check{P}Tf & & \uparrow \check{P}f \\
 (PS')^n & \xrightarrow{\lambda_{S'}} & PTS' & \xrightarrow{\check{P}\sigma'} & PS'
 \end{array} \tag{50}$$

Now consider the following calculation:

$$\begin{aligned}
 \llbracket \heartsuit_\lambda(\dots, \psi_i, \dots) \rrbracket^{\mathbb{S}} &= (\check{P}\sigma)\lambda_S(\dots, \llbracket \psi_i \rrbracket^{\mathbb{S}}, \dots) && \text{(semantics } \heartsuit_\lambda \text{ in } \mathbb{S}, \text{ as in (46))} \\
 &= (\check{P}\sigma)\lambda_S(\dots, (\check{P}f)\llbracket \psi_i \rrbracket^{\mathbb{S}'}, \dots) && \text{(induction hypothesis)} \\
 &= (\check{P}\sigma)(\check{P}Tf)\lambda_{S'}(\dots, \llbracket \psi_i \rrbracket^{\mathbb{S}'}, \dots) && \text{(naturality of } \lambda, \text{ (50) left)} \\
 &= (\check{P}f)(\check{P}\sigma)\lambda_{S'}(\dots, \llbracket \psi_i \rrbracket^{\mathbb{S}'}, \dots) && \text{(} f \text{ is a morphism, (50) right)} \\
 &= (\check{P}f)\llbracket \heartsuit_\lambda(\dots, \psi_i, \dots) \rrbracket^{\mathbb{S}'} && \text{(semantics } \heartsuit_\lambda \text{ in } \mathbb{S}', \text{ as in (46))}
 \end{aligned}$$

showing that (49) holds for $\varphi = \heartsuit_\lambda(\psi_0, \dots, \psi_{n-1})$ indeed. QED

Concerning the property of expressiveness, we may obtain a general result if we put some constraints on the signature Λ .

Definition 6.13 Let Λ be a modal signature for a set functor T . We say that Λ is *separating* for T , if for all sets S and all pairs of distinct objects $\sigma_0, \sigma_1 \in TS$ there is a $\lambda \in \Lambda$ and a tuple $(A_0, \dots, A_{n-1})$ such that *exactly* one of the two objects σ_i belongs to the set $\lambda_S(A_0, \dots, A_{n-1})$. $\triangleleft$

Example 6.14 (1) The box relation lifting $\lambda^\square$ is separating on its own, for the powerset functor P . To see this, consider two subsets $X, Y \in PS$. If X and Y are distinct, suppose without loss of generality that $Y \not\subseteq X$, so that $Y \notin \lambda^\square(X) = \{U \in PS \mid U \subseteq X\}$.

(2) The predicate liftings associated with the graded modalities are jointly separating. To see this, consider two bags $\beta_0, \beta_1 : S \rightarrow \mathbb{N}^\infty$ over some set S . If β_0 and β_1 are distinct, then we must have $\beta_0(s) \neq \beta_1(s)$, for some $s \in S$. Without loss of generality we may assume that $\beta_0(s) < \beta_1(s)$, so that in particular, $\beta_0(s)$ belongs to $\mathbb{N}$, say, $\beta_0(s) = m$. Recall that $\lambda_S^{m+1}(\{s\}) = \{\beta \in BS \mid \beta(s) \geq m+1\}$, so that we find $\beta_0 \notin \lambda_S^{m+1}(\{s\})$ but $\beta_1 \in \lambda_S^{m+1}(\{s\})$.

Theorem 6.15 Let Λ be a separating modal signature for the set functor T . Then the language ML_Λ is expressive on the class of image-finite T -coalgebras: Given any two pointed T_ω -models $(\mathbb{S}, s)$ and $(\mathbb{S}', s')$ we have

$$(\mathbb{S}, s) \equiv_\Lambda (\mathbb{S}', s') \text{ implies } (\mathbb{S}, s) \simeq_T (\mathbb{S}', s'). \tag{51}$$

Proof. For notational simplicity we will confine ourselves to a setting where all predicate liftings in Λ are unary, leaving the (routine) generalisation to an arbitrary signature as an

exercise. It will also be convenient to assume that T preserves inclusions. Furthermore, we will only treat the special case where the coalgebras $\mathbb{S}$ and $\mathbb{S}'$ coincide (i.e., $\mathbb{S} = \mathbb{S}'$); the general case, where the two coalgebras are distinct, can easily be reduced to this by considering their disjoint union.

So let Λ be a separating set of unary predicate liftings for some set functor T , and let $\mathbb{S} = (S, \sigma, V)$ be an image-finite T -model, that is, $\sigma : S \rightarrow T_\omega S$. To avoid notational clutter we will simply write $\equiv$ for the equivalence relation $\equiv_\Lambda$.

Our aim is prove that $s_0 \equiv s_1$ implies $s_0 \simeq s_1$, for all $s_0, s_1 \in S$. Clearly then it suffices to show that the relation $\equiv$ is contained in the kernel of some coalgebra morphism. We will show that in fact we may define coalgebra structure $\bar{\sigma} : \bar{S} \rightarrow T\bar{S}$ on the set $\bar{S}$ of $\equiv$ -cells in such a way that the quotient map $q : S \rightarrow \bar{S}$ becomes a coalgebra morphism:

$$\begin{array}{ccc} S & \xrightarrow{q} & \bar{S} \\ \sigma \downarrow & & \downarrow \bar{\sigma} \\ TS & \xrightarrow{Tq} & T\bar{S} \end{array}$$

Now suppose that we can show that

$$q(s_0) = q(s_1) \text{ implies } (Tq)(\sigma s_0) = (Tq)(\sigma s_1), \quad (52)$$

then putting

$$\bar{\sigma}(\bar{s}) := (Tq)(\sigma s)$$

would give a correctly defined map, for which the quotient map q is trivially a coalgebra morphism.

That is, we have reduced our problem to finding a proof for (52), and to this aim we reason by contraposition: Assuming that

$$(Tq)(\sigma s_0) \neq (Tq)(\sigma s_1), \quad (53)$$

we will show that $q(s_0) \neq q(s_1)$ (that is, $s_0 \not\equiv_\Lambda s_1$). It follows by separation from (53) that there is some $\lambda \in \Lambda$ and some $A \subseteq \bar{S}$ such that (without loss of generality) we have

$$(Tq)(\sigma s_0) \in \lambda_{\bar{S}}(A), \text{ but } (Tq)(\sigma s_1) \notin \lambda_{\bar{S}}(A). \quad (54)$$

Our purpose is now is find a formula witnessing this, in the sense that this formula holds at s_0 but not at s_1 : this would show indeed that $s_0 \not\equiv s_1$, and so $q(s_0) \neq q(s_1)$. Note that $A \subseteq \bar{S}$ simply means that A is a collection of equivalence classes.

Now if we would have $\bigcup A = \llbracket \varphi \rrbracket$ we would be done immediately since in this case we could prove that $\mathbb{S}, s_0 \models \heartsuit_\lambda \varphi$ while $\mathbb{S}, s_1 \not\models \heartsuit_\lambda \varphi$. To see this, note that we have $s_i \models \heartsuit_\lambda \varphi$ iff $\sigma(s_i) \in \lambda_S(\llbracket \varphi \rrbracket)$ and that $\bigcup A = \llbracket \varphi \rrbracket$ implies $\lambda_S(\llbracket \varphi \rrbracket) = \lambda_S(\bigcup A) = \lambda_S(\check{P}q(A)) = (\check{P}Tq)(\lambda_{\bar{S}}(A))$, where the last identity is by naturality of λ . It is thus immediate by (54) that $\sigma s_0 \in \lambda_S(\llbracket \varphi \rrbracket)$ but $\sigma s_1 \notin \lambda_S(\llbracket \varphi \rrbracket)$, as required.

In the general case we need to work harder. But since $\mathbb{S}$ is a T_ω -coalgebra, and T preserves inclusions, there is a *finite* subset $X \subseteq S$ such that both σs_0 and σs_1 belong to the set TX .

We leave it for the reader to verify that there is a formula $\varphi \in \text{ML}_\Lambda$ that characterizes, *within* X , the union $\bigcup A$, in the sense that

$$\text{for all } x \in X : \quad \mathbb{S}, x \Vdash \varphi \text{ iff } x \in \bigcup A,$$

or equivalently, since $x \in \bigcup A$ is another way of saying that $q(x) \in A$:

$$X \cap \llbracket \varphi \rrbracket = X \cap (\check{P}q)A. \quad (55)$$

We now claim that, for this formula φ , we have

$$\mathbb{S}, s_0 \Vdash \heartsuit_\lambda \varphi \text{ but } \mathbb{S}, s_1 \not\Vdash \heartsuit_\lambda \varphi. \quad (56)$$

To prove this, we first observe that by the semantics of $\heartsuit_\lambda$ we have that $\mathbb{S}, s_i \Vdash \heartsuit_\lambda \varphi$ iff $\sigma(s_i) \in \lambda_S(\llbracket \varphi \rrbracket)$, while it follows from (54) that $\sigma(s_0) \in (\check{P}Tq)\lambda_{\bar{S}}(A)$ but $\sigma(s_1) \notin (\check{P}Tq)\lambda_{\bar{S}}(A)$. Hence, because both $\sigma(s_0)$ and $\sigma(s_1)$ belong to TX , it suffices to show that

$$TX \cap (\check{P}Tq)\lambda_{\bar{S}}(A) = TX \cap \lambda_S \llbracket \varphi \rrbracket. \quad (57)$$

We will establish this by chasing the diagram below, where we use a trick to interpret the intersection with X in (55) and with TX in (57) using the inclusion map $\iota : X \hookrightarrow S$. That is, we observe that $\check{P}\iota : PS \rightarrow PX$ is given by $U \mapsto X \cap U$; as a consequence another way of formulating (55) is:

$$(\check{P}\iota)(\llbracket \varphi \rrbracket) = (\check{P}\iota)(\check{P}q)(A). \quad (58)$$

Similarly, since T preserves inclusions we have that $T\iota : TX \rightarrow TS$ is the inclusion map witnessing that $TX \subseteq TS$, and so $\check{P}T\iota : PTS \rightarrow PTX$ is given by $\Sigma \mapsto (TX) \cap \Sigma$.

$$\begin{array}{ccc} \bar{S} & PS \xrightarrow{\lambda_{\bar{S}}} PT\bar{S} & TX \cap (\check{P}Tq)\lambda_{\bar{S}}(A) \\ q \uparrow & \check{P}q \downarrow & = (\check{P}T\iota)(\check{P}Tq)\lambda_{\bar{S}}(A) \quad (\text{just discussed}) \\ S & PS \xrightarrow{\lambda_S} PTS & = \lambda_X(\check{P}\iota)(\check{P}q)(A) \quad (\text{naturality of } \lambda) \\ \iota \uparrow & \check{P}\iota \downarrow & = \lambda_X(\check{P}\iota)(\llbracket \varphi \rrbracket) \quad (\text{see (58)}) \\ X & PX \xrightarrow{\lambda_X} PTX & = (\check{P}T\iota)\lambda_S(\llbracket \varphi \rrbracket) \quad (\text{naturality of } \lambda) \\ & & = TX \cap \lambda_S \llbracket \varphi \rrbracket \quad (\text{just discussed}) \end{array}$$

This proves (57), and therefore (56). That is, we have shown that $s_0 \not\equiv s_1$, on the assumption that $(Tq)(\sigma s_0) \neq (Tq)(\sigma s_1)$. This means that (52) holds, and as we argued already, this suffices to prove the Theorem. QED $\bullet$

6.4 Finite model property

In this subsection we will show that the coalgebraic modal logic ML_Λ has the (strong) finite model property. That is, we will show that any satisfiable ML_Λ -formula φ can in fact be satisfied in a *finite* coalgebra of which the size (number of states) is exponentially bounded by the size of φ . We will prove this result by adapting the method of filtration, which is well known in the theory of standard modal logic, to the more general coalgebraic setting.

First we need some preliminary definitions.

Definition 6.16 The collection $\mathbf{Sfor}(\varphi)$ of *subformulas* of a ML_Λ -formula φ is defined in the standard way. The *size* $|\varphi|$ of a φ is defined as its number of subformulas: $|\varphi| := |\mathbf{Sfor}(\varphi)|$.

A set of formulas Σ is called *subformula-closed* if it is closed under taking subformulas, that is, if $\mathbf{Sfor}(\varphi) \subseteq \Sigma$ for all $\varphi \in \Sigma$. $\triangleleft$

The idea behind the method of filtration is fairly simple: given a subformula-closed set Σ and a T -model $\mathbb{S} = (S, \sigma, V)$, define on S a suitable equivalence relation $\equiv_\Sigma$ of finite index, build a new T -model $\bar{\mathbb{S}}$ on the finite set of $\equiv_\Sigma$ -cells, and show that any state $s \in S$ satisfies the same formulas in $\mathbb{S}$ as its cell $\bar{s}$ does in the filtrated model $\bar{\mathbb{S}}$.

Definition 6.17 Let Σ be a finite, subformula closed set of formulas in $\text{ML}_\Lambda(\mathbf{Q})$, and let $\mathbb{S} = (S, \sigma, V)$ be a T -model. We define $\equiv_\Sigma \subseteq S \times S$ as the equivalence relation given by

$$s \equiv_\Sigma t \text{ iff for all } \varphi \in \Sigma : \mathbb{S}, s \Vdash \varphi \iff \mathbb{S}, t \Vdash \varphi,$$

and denote the $\equiv_\Sigma$ -cell of a state s as $\bar{s}$. We also let $\bar{S} := \{\bar{s} \mid s \in S\}$ denote the set of cells, and let $q : s \mapsto \bar{s}$ denote the quotient map $q : S \rightarrow \bar{S}$. $\triangleleft$

In order to define a coalgebra map $\bar{\sigma} : \bar{S} \rightarrow T\bar{S}$, we may pick any *choice function* $c : \bar{S} \rightarrow S$, and define $\bar{\sigma} := Tq \circ \sigma \circ c$, cf. the diagram below:

$$\begin{array}{ccc} S & \xleftarrow{c} & \bar{S} \\ \downarrow \sigma & \searrow q & \downarrow \bar{\sigma} \\ TS & \xrightarrow{Tq} & T\bar{S} \end{array}$$

Here we call $c : \bar{S} \rightarrow S$ a choice function if c picks an element from each $\equiv_\Sigma$ -cell; in other words, we require that $q \circ c = \text{id}_{\bar{S}}$ and $\ker(c \circ q) \subseteq \equiv_\Sigma$.

Note that while the ‘outer’ rectangle of the above diagram commutes by definition, the ‘inner’ one need not commute: it will generally not be possible to define a coalgebra map on $\bar{S}$ for which the quotient map q is a coalgebra morphism. Fortunately, for our purposes we don’t need such a coalgebra morphism.

Definition 6.18 Let Σ be a finite, subformula closed set of formulas in $\text{ML}_\Lambda(\mathbf{Q})$, and let $\mathbb{S} = (S, \sigma, V)$ be a T -model. A Σ -filtration of $\mathbb{S}$ is any T -model $\bar{\mathbb{S}} = (\bar{S}, \bar{\sigma}, \bar{V})$ such that

- (1) $\bar{S} = S/\equiv_\Sigma$ is the class of $\equiv_\Sigma$ -cells,
- (2) $\bar{\sigma} = Tq \circ \sigma \circ c$ for some choice function $c : \bar{S} \rightarrow S$, and
- (3) $\bar{V}(q) = \{\bar{s} \mid s \in V(q)\}$ for $q \in \Sigma \cap \mathbf{Q}$. $\triangleleft$

Observe that filtrations are not unique, they depend in particular on the choice of the choice function c .

We can now prove the following Filtration Lemma.

Theorem 6.19 (Filtration Lemma) *Let Λ be a modal signature for a set functor T , and let $\Sigma \subseteq \text{ML}_\Lambda$ be a finite subformula-closed set of formulas. Furthermore, let $\mathbb{S} = (S, \sigma, V)$ be a T -model, and let $\bar{\mathbb{S}} = (\bar{S}, \bar{\sigma}, \bar{V})$ be a Σ -filtration of $\mathbb{S}$. Then for all formulas $\varphi \in \Sigma$ we have*

$$\mathbb{S}, s \Vdash \varphi \text{ iff } \bar{\mathbb{S}}, \bar{s} \Vdash \varphi, \text{ for all states } s \in S. \quad (59)$$

Proof. First, note that the statement (59) is equivalent to

$$\llbracket \varphi \rrbracket^{\mathbb{S}} = (\check{P}q)\llbracket \varphi \rrbracket^{\bar{\mathbb{S}}}. \quad (60)$$

This means that we may use (60) as the inductive hypothesis in our inductive proof.

Second, since we have $s \equiv c(\bar{s})$ for all $s \in S$ the statement (59) is also equivalent to

$$\llbracket \varphi \rrbracket^{\bar{\mathbb{S}}} = (\check{P}c)\llbracket \varphi \rrbracket^{\mathbb{S}}, \quad (61)$$

as long as we restrict to formulas $\varphi \in \Sigma$. Since we only want to establish (59) for formulas in Σ , it therefore suffices to prove (61). We will do so by a straightforward formula induction.

Turning to the actual inductive proof of (61), we only consider the inductive step where $\varphi = \heartsuit_\lambda(\psi_0, \dots, \psi_{n-1})$, and for notational simplicity we confine ourselves to the case where λ is unary, i.e. $\varphi = \heartsuit_\lambda\psi$ for some formula ψ to which the inductive hypothesis applies. Now consider the following diagram:

$$\begin{array}{ccccc} P\bar{S} & \xrightarrow{\lambda_{\bar{S}}} & PT\bar{S} & \xrightarrow{\check{P}\bar{\sigma}} & P\bar{S} \\ \check{P}q \downarrow & & \check{P}Tq \downarrow & & \uparrow \check{P}c \\ PS & \xrightarrow{\lambda_S} & PTS & \xrightarrow{\check{P}\sigma} & PS \end{array} \quad (62)$$

$$\begin{aligned} \llbracket \heartsuit_\lambda\psi \rrbracket^{\bar{\mathbb{S}}} &= (\check{P}\bar{\sigma})\lambda_{\bar{S}}(\llbracket \psi \rrbracket^{\bar{\mathbb{S}}}) && \text{(semantics of } \heartsuit_\lambda) \\ &= (\check{P}(Tq \circ \sigma \circ c))\lambda_{\bar{S}}(\llbracket \psi \rrbracket^{\bar{\mathbb{S}}}) && \text{(definition } \bar{\sigma}) \\ &= (\check{P}c)(\check{P}\sigma)(\check{P}Tq)\lambda_{\bar{S}}(\llbracket \psi \rrbracket^{\bar{\mathbb{S}}}) && \text{(functoriality)} \\ &= (\check{P}c)(\check{P}\sigma)\lambda_S(\check{P}q)(\llbracket \psi \rrbracket^{\bar{\mathbb{S}}}) && \text{(naturality of } \lambda) \\ &= (\check{P}c)(\check{P}\sigma)\lambda_S(\llbracket \psi \rrbracket^{\mathbb{S}}) && \text{(induction hypothesis (60))} \\ &= (\check{P}c)\llbracket \heartsuit_\lambda\psi \rrbracket^{\mathbb{S}} && \text{(semantics of } \heartsuit_\lambda) \end{aligned}$$

In other words, we have established (61) for $\varphi = \heartsuit_\lambda\psi$, as required. QED

Corollary 6.20 (Strong Finite Model Property) *Let φ be a formula in ML_Λ , where Λ is a modal signature for a set functor T . If φ is satisfiable in some T -model, then it is satisfiable in a finite T -model (S, σ, V) such that $|S| \leq 2^{|\varphi|}$.*

Proof. Fix an ML_Λ -formula φ , and let $\Sigma := \text{Sfor}(\varphi)$. It follows by the filtration lemma that φ , if satisfiable in some pointed T -model $(\mathbb{S}, s)$, also holds at the state $\bar{s}$, in any filtration $\bar{\mathbb{S}}$ of $\mathbb{S}$. This proves the theorem, since it easily follows from the definition of the relation $\equiv_\Sigma$ that $|\bar{S}| \leq |P\Sigma| = 2^{|\varphi|}$. QED

6.5 Predicate liftings as coalgebra type changers

This short section presents a slightly different perspective in which predicate liftings provide natural ways to transform T -coalgebras to neighbourhood frames. We first consider a simplified version, in which T -coalgebras are transformed to Kripke frames.

Example 6.21 Define a *natural relation* for T to be a natural transformation $\mu : T \rightarrowtail P$.

Given such a natural relation $\mu : T \rightarrowtail P$, we can transform a T -coalgebra $\mathbb{S} = (S, \sigma)$ into a Kripke frame $\mathbb{S}^\mu := (S, \mu_S \circ \sigma)$. By naturality of μ , any T -homomorphism $f : \mathbb{S} \rightarrow \mathbb{S}'$ is also a bounded morphism $f : \mathbb{S}^\mu \rightarrow (\mathbb{S}')^\mu$. To check this, one may easily verify that $Pf \circ (\mu_S \circ \sigma) = (\mu_{S'} \circ \sigma') \circ f$ by chasing the diagram below:

$$\begin{array}{ccccc} S & \xrightarrow{\sigma} & TS & \xrightarrow{\mu_S} & PS \\ f \downarrow & & \downarrow Tf & & \downarrow Pf \\ S' & \xrightarrow{\sigma'} & TS' & \xrightarrow{\mu_{S'}} & PS' \end{array}$$

Connecting this to logic, with any natural relation μ we may associate a modality $\langle \mu \rangle$ for T -coalgebras, with the following interpretation:

$$\mathbb{S}, s \Vdash \langle \mu \rangle \varphi \text{ iff } \llbracket \varphi \rrbracket^{\mathbb{S}} \cap \mu_S \sigma(s) \neq \emptyset.$$

As an example, recall from Fact 5.16(2) that for smooth and standard functors T , we have a natural transformation $Base^T : T_\omega \rightarrowtail P_\omega$. Hence, we may take the *Base* operation as a way to transform T_ω -coalgebras into P_ω -coalgebras, that is, image-finite Kripke frames. Naturality ensures that every morphism between T -coalgebras is also a bounded morphism between the underlying Kripke frames.

As we will see now, predicate liftings can be seen as generalisations of this phenomenon, where we move from Kripke frames to the more general setting of neighbourhood frames. For this general setting we introduce the *transpose* of a predicate lifting, cf. Definition A.4. This notion is based on the correspondence between maps $A \rightarrow PB$ and maps $B \rightarrow PA$ — we have seen this correspondence already in the coalgebraic presentation of a valuation $V : \mathbf{Q} \rightarrow PS$ as a colouring $V^\flat : S \rightarrow PQ$.

Definition 6.22 Given a map $\alpha : A \rightarrow PB$ we define its *transposed map* $\alpha^\flat : B \rightarrow PA$ by putting $\alpha^\flat(b) := \{a \in A \mid b \in \alpha(a)\}$.

Extending this definition, given an n -ary predicate lifting for the set functor T , we define its *transpose* $\lambda^\flat$ as the set-indexed family of maps

$$\lambda_S^\flat : TS \rightarrow P(P^n(S))$$

given by $\lambda_S^\flat(\sigma) := (\lambda_S)^\flat(\sigma)$. $\triangleleft$

By the naturality of predicate liftings we obtain the following proposition, which shows that predicate liftings indeed generalise the natural relations of Example 6.21.

Proposition 6.23 *If $\lambda : P^n \rightarrow PT$ then $\lambda^\flat$ is a natural transformation*

$$\lambda^\flat : T \rightarrow \check{P} \circ \check{P}^n.$$

It follows from this proposition that any predicate lifting λ induces a transformation of T -coalgebras to n -ary neighbourhood frames. We confine attention to the unary case.

Definition 6.24 Let $\lambda : \check{P} \rightarrow \check{P}T$ be a unary predicate lifting for the set functor T . Given a T -coalgebra $\mathbb{S} = (S, \sigma)$, we let $\mathbb{S}^\lambda$ denote the neighbourhood frame $\mathbb{S}^\lambda := (S, \lambda^\flat \circ \sigma)$; given a function $f : S \rightarrow S'$ we define $f^\lambda := f$. $\triangleleft$

The following proposition is easy to verify.

Proposition 6.25 *Let $\lambda : \check{P} \rightarrow \check{P}T$ be a unary predicate lifting for the set functor T . Then the construction $(\cdot)^\lambda$ is a functor from the category of T -coalgebras to the category of neighbourhood frames (N -coalgebras).*

The following proposition, which is easy to verify, provides a slightly different perspective on the concept of separation.

Proposition 6.26 *Let Λ be a modal signature for a set functor T . Then Λ is separating iff, for every set S , the collection $(\lambda_S^\flat)_{\lambda \in \Lambda}$ of transposed functions is jointly injective (i.e., for any pair of distinct objects $\tau_0, \tau_1 \in TS$ there is a $\lambda \in \Lambda$ such that $\lambda_S^\flat(\tau_0) \neq \lambda_S^\flat(\tau_1)$).*

6.6 Predicate liftings and the Yoneda lemma

In the final subsection of this chapter we take a slightly different perspective on the modalities that are given by predicate liftings, seeing them as describing certain *admissible patterns*. This perspective will also reveal *how many* predicate liftings of each different arity there are.

The key observation here is that there is a natural bijection between PS (subsets of S) and 2^S (characteristic functions on S).

Definition 6.27 Given a subset $X \subseteq S$, we define the *characteristic function of X* as the map $\chi_X^S : S \rightarrow 2$ given by

$$\chi_X^S(s) := \begin{cases} 1 & \text{if } s \in X \\ 0 & \text{if } s \notin X. \end{cases}$$

In case S is understood, we may drop the superscript ‘ S ’. Conversely, given a map $\chi : S \rightarrow 2$, we shall call $\chi^{-1}(1) \in PS$ the *subset determined by χ* . $\triangleleft$

Remark 6.28 This correspondence reaches far enough for us to think of the contravariant power set functor as the functor 2^- that associates with a set S the collection 2^S of functions from S to $2 = \{0, 1\}$, and with an arrow $f : S' \rightarrow S$ the ‘precomposition map’ that assigns to an arbitrary characteristic function $\chi : S \rightarrow 2$ the function $\chi \circ f : S' \rightarrow 2$.

From this point of view, a unary predicate lifting is a way of transforming arrows $S \rightarrow 2$ into arrows $TS \rightarrow 2$. In particular, any arrow $\gamma : T2 \rightarrow 2$ (that is, any characteristic function χ_Γ corresponding to a subset $\Gamma \subseteq T2$), induces a unary predicate lifting: Given an arrow $\chi : S \rightarrow 2$, simply consider the arrow $\gamma \circ T\chi$, as in the diagram below:

$$\begin{array}{ccc} S & \xrightarrow{\chi} & 2 \\ & \searrow \scriptstyle \gamma \circ T\chi & \\ TS & \xrightarrow{T\chi} T2 & \xrightarrow{\gamma} 2 \end{array}$$

Formulated in terms of subsets rather than characteristic functions, we arrive at the following definition.

Definition 6.29 Given an object $\Gamma \subseteq T2$, let $\hat{\Gamma}$ be the following set-indexed family of operations. For a set S , we define

$$\hat{\Gamma}_S : PS \rightarrow PTS,$$

by putting, for any $X \subseteq S$:

$$\hat{\Gamma}_S(X) := \{\sigma \in TS \mid (T\chi_X^S)(\sigma) \in \Gamma\},$$

where $\chi_X^S : S \rightarrow 2$ is the characteristics map associated with X . Where $\hat{\Gamma}$ is a predicate lifting, we shall denote its associated modality as $\heartsuit_\Gamma$ rather than as $\heartsuit_{\hat{\Gamma}}$. $\triangleleft$

Remark 6.30 Taking a glance at the *modalities* that are induced by subsets of the set $T2$, we consider the interpretation of the formula $\heartsuit_\Gamma \varphi$ in the T -model $\mathbb{S} = (S, \sigma, V)$. If we represent the set $\llbracket \varphi \rrbracket$ with its characteristic function $\chi_{\llbracket \varphi \rrbracket}^S$, applying the functor T to this arrow we obtain

$$TS \xrightarrow{T\chi_{\llbracket \varphi \rrbracket}^S} T2.$$

Now think of the elements of $T2$ as ‘ T -patterns’, then the above arrow associates a T -pattern with each object $\tau \in TS$. We can say that the formula $\heartsuit_\Gamma \varphi$ holds at s if the pattern $(T\chi_{\llbracket \varphi \rrbracket}^S)(\sigma(s))$ associated with $\sigma(s) \in TS$ is *admissible*, i.e., belongs to the set Γ , or, equivalently, that $(\chi_\Gamma^{T2} \circ T\chi_{\llbracket \varphi \rrbracket}^S)(\sigma(s)) = 1$.

Example 6.31 As an example, consider the binary tree functor $Id \times Id$. The set $T2$ consists of four patterns: $T2 = \{(0,0), (0,1), (1,0), (1,1)\}$. As an example, take the set $\Gamma = \{(0,0), (1,1)\}$; it is not hard to see that it induces the predicate lifting $\hat{\Gamma} : \check{P} \rightarrow \check{P} \circ (Id \times Id)$ given by

$$\hat{\Gamma}_S(X) := X \times X \cup (S \setminus X) \times (S \setminus X).$$

The associated modality $\heartsuit_\Gamma$ has the following semantics:

$\mathbb{S}, s \Vdash \heartsuit_\Gamma \varphi$ iff the two successors of s either both satisfy or both falsify φ .

Proposition 6.32 For any $\Gamma \subseteq T2$, the collection $\hat{\Gamma}$ of maps constitutes a predicate lifting $\hat{\Gamma} : \check{P} \rightarrow \check{P}T$.

Proof. Given a map $f : S' \rightarrow S$, we need to check that the following diagram commutes:

$$\begin{array}{ccccc}
 S' & & PS' & \xrightarrow{\hat{\Gamma}_{S'}} & PTS' \\
 f \downarrow & & \check{P}f \uparrow & & \uparrow \check{P}Tf \\
 S & & PS & \xrightarrow{\hat{\Gamma}_S} & PTS
 \end{array}$$

This follows by the following chain of identities, for an arbitrary subset $X \subseteq S$:

$$\begin{aligned}
 (\hat{\Gamma}_{S'} \circ \check{P}f)(X) &= \hat{\Gamma}_{S'}(\check{P}f(X)) && \text{(obvious)} \\
 &= \{\sigma' \in TS' \mid (T\chi_{\check{P}f(X)}^{S'})(\sigma') \in \Gamma\} && \text{(definition } \hat{\Gamma}_{S'}) \\
 &= \{\sigma' \in TS' \mid (T(\chi_X^S \circ f))(\sigma') \in \Gamma\} && (*) \\
 &= \{\sigma' \in TS' \mid (T\chi_X^S)((Tf)(\sigma')) \in \Gamma\} && \text{(functoriality of } T) \\
 &= \{\sigma' \in TS' \mid (Tf)(\sigma') \in \hat{\Gamma}_S(X)\} && \text{(definition } \hat{\Gamma}_S) \\
 &= (\check{P}Tf)(\hat{\Gamma}_S(X)) && \text{(definition } \check{P}Tf) \\
 &= (\check{P}Tf \circ \hat{\Gamma}_S)(X) && \text{(obvious)}
 \end{aligned}$$

Here the identity $(*)$ is immediate by the observation that

$$\chi_X^S \circ f = \chi_{\check{P}f(X)}^{S'}$$

as is revealed by a straightforward verification, for an arbitrary element $s' \in S'$: $\chi_X^S \circ f(s') = 1$ iff $f(s') \in X$ iff $s' \in f^{-1}(X) = (\check{P}f)(X)$ iff $\chi_{\check{P}f(X)}^{S'}(s') = 1$. QED

Interestingly, we may prove that *all* unary predicate liftings are of this form, and this result generalises to predicate liftings of arbitrary arity. This is the main content of the following theorem, which is in fact the instantiation of the well-known *Yoneda Lemma* to the setting of predicate liftings.

Theorem 6.33 *For any set functor T there is a natural bijection between the set of n -ary predicate liftings for T and the power set of $T(Pn)$.*

The key observations underlying the proof of this theorem are the natural correspondences

$$(PS)^n \sim (2^S)^n \sim (2^n)^S \sim (Pn)^S$$

between n -tuples of subsets of S , n -tuples of characteristic functions, maps from S to 2^n , and maps from S to Pn .

Proof. We confine the proof of this result to providing an n -ary predicate lifting for each subset of TPn , and vice versa.

So let $\Gamma \subseteq T(Pn)$ be a set of admissible n -ary T -patterns. The associated predicate lifting $\hat{\Gamma} : P \rightarrow PT$ is obtained by a straightforward generalisation of Definition 6.29. For the

details, take an arbitrary tuple $X = (X_0, \dots, X_{n-1}) \in (PS)^n$ of subsets of S , and represent this tuple as the map $\chi_X := (\chi_0, \dots, \chi_{n-1})$ of associated characteristic functions, then we have $\chi_X : S \rightarrow 2^n$, and so $T\chi_X : TS \rightarrow T(2^n)$. Now put

$$\widehat{\Gamma}_S(X_0, \dots, X_{n-1}) := \{\tau \in TS \mid (T\chi_X)(\tau) \in \Gamma\}.$$

We leave it for the reader to verify that this collection of maps indeed provides a natural transformation $\widehat{\Gamma} : \check{P}^n \rightarrow \check{P}T$.

For the opposite direction, let $\lambda : \check{P}^n \rightarrow \check{P}T$ be an n -ary predicate lifting. Our aim is to find a subset $\Theta_\lambda \subseteq T(2^n)$ such that $\lambda = \widehat{\Theta_\lambda}$.

The easiest way to proceed is by thinking of λ_S as a way to transform arrows $S \rightarrow 2^n$ into arrows $TS \rightarrow 2$, that is, $\lambda_S : (2^n)^S \rightarrow 2^{TS}$. To find the object Θ_λ , we take a *special* set S , viz., the set 2^n itself, and, as input for λ , a *special* arrow, viz., the *identity* arrow $\text{id}_{2^n} : 2^n \rightarrow 2^n$. Then we let $\Theta_\lambda \subseteq T(2^n)$ be the subset of $T(2^n)$ that is determined by the image of this arrow id_{2^n} under λ .

If we think of λ as a set-indexed family of maps $\lambda_S : (PS)^n \rightarrow PTS$, we may define

$$\Theta_\lambda := \lambda(\mathcal{U}),$$

where $\mathcal{U}$ is the distinguished element of the set $P^n(Pn)$ corresponding to the identity map id_{2^n} , that is, $\mathcal{U} = (\mathcal{U}_0, \dots, \mathcal{U}_{n-1})$ with

$$\mathcal{U}_i := \{U \subseteq n \mid i \in U\}.$$

As mentioned, we leave it for the reader to verify that the maps just defined are each other's inverse, i.e., that $\lambda = \widehat{\Theta_\lambda}$ for all predicate liftings λ , and that $\Gamma = \Theta_{\widehat{\Gamma}}$ for all $\Gamma \in PTn$, some $n \in \omega$. QED

7 One-step logic

In this chapter we will zoom in on a micro-version of coalgebraic modal logic that we call *one-step logic*. For this purpose we introduce *one-step models*; intuitively, these are *windows* over a T -model, only allowing access to the coalgebraic unfolding of one single state. Formally one-step models are simply defined as triples (S, σ, m) with $\sigma \in TS$ and m a marking of states in S with a set of variables. In particular, no coalgebra map is assumed.

Then, given a modal signature Λ , we will introduce a very simple modal language $1ML_\Lambda$ for describing properties of one-step models, the *one-step language* associated with Λ . The so-called *one-step formulas* of $1ML_\Lambda$ will be based on a set A of propositional variables, which should be seen as place-holders for proper ML_Λ -formulas. In particular, they are *different* from the proposition letters of the languages — to avoid confusion the latter could be encoded as (nullary) predicate liftings, see Remark 6.11. Characteristic of one-step formulas is the syntactic restriction that all occurrences of variables are in the scope of exactly one modality; thus the one-step formalism can be seen as a very simple fragment of the full modal language.

The point of focussing on this ‘local’, one-step version of coalgebraic modal logic is that many properties of coalgebraic modal logic are in fact already determined at this one-step level. As we will see in this chapter and further on, this applies both to model-theoretic properties and to derivation systems.

Throughout this chapter we fix a set functor T and a modal signature Λ .

7.1 One-step syntax and semantics

We can now turn to the formal definitions of one-step logic. We start with syntax. Given any set E of expressions, we let $PL(E)$ denote the set of (*boolean*) *propositional* or *rank-0* formulas over E , given by the following grammar:

$$\pi ::= e \in E \mid \perp \mid \top \mid \pi_0 \vee \pi_1 \mid \pi_0 \wedge \pi_1 \mid \neg \pi.$$

Definition 7.1 Fix a set Λ of predicate liftings, and a set A of propositional variables. Given any set Φ of formulas, we define

$$\Lambda(\Phi) := \{\heartsuit_\lambda(\varphi_0, \dots, \varphi_{n-1} \mid \lambda \in \Lambda, \text{ar}(\lambda) = n, \varphi_0, \dots, \varphi_{n-1} \in \Phi\}.$$

We shall call

$$1ML_\Lambda(A) := PL(\Lambda(PL(A))).$$

the *one-step language* for Λ over A , and refer to formulas $\alpha, \beta \in 1ML_\Lambda$ as *one-step* or *rank-1* Λ -formulas. For both $PL(A)$ and $1ML_\Lambda(A)$ we may use standard propositional abbreviations such as the connectives $\rightarrow$ and $\leftrightarrow$. $\triangleleft$

A more direct characterisation of the language $1ML_\Lambda(A)$ is that it consists of those $ML_\Lambda(A)$ -formulas in which every variable $a \in A$ occurs in the scope of *exactly one modality*. As yet another alternative, the set $1ML_\Lambda(A)$ can be characterised by the following grammar:

$$\alpha ::= \heartsuit_\lambda(\pi_0, \dots, \pi_{n-1}) \mid \perp \mid \top \mid \alpha_0 \vee \alpha_1 \mid \alpha_0 \wedge \alpha_1 \mid \neg \alpha,$$

where $\lambda \in \Lambda$ and $\pi_i \in PL(A)$.

Example 7.2 (1) With $A = \{a, b, c\}$ and $\Lambda = \{\Box, \Diamond\}$, examples of one-step formulas are $\Diamond(a \wedge b)$, $(\Box((a \vee a) \wedge \perp) \vee \top$. The formulas $\Diamond\Diamond a$ and $a \vee \Box\perp$ are not one-step formulas.

One-step formulas are naturally interpreted in *one-step models*, which consist of a one-step frame together with a marking. As mentioned, a one-step T -model can be seen as a window over a T -model, or as a potential unfolding of one state in a T -model.

Definition 7.3 A *one-step T -frame* is a pair (S, σ) with $\sigma \in TS$. A *one-step T -model* over a set A of variables is a triple (S, σ, m) such that (S, σ) is a one-step T -frame and $m : S \rightarrow PA$ is a A -marking on S . $\triangleleft$

Definition 7.4 Given a marking $m : S \rightarrow PA$, we define the *0-step interpretation* $\llbracket \pi \rrbracket_m^0 \subseteq S$ of $\pi \in \text{PL}(A)$ by the obvious induction:

$$\begin{aligned} \llbracket a \rrbracket_m^0 &:= m^b(a) \quad \left(= \{s \in S \mid a \in m(s)\} \right) \\ \llbracket \top \rrbracket_m^0 &:= S \\ \llbracket \perp \rrbracket_m^0 &:= \emptyset \\ \llbracket \neg\pi \rrbracket_m^0 &:= S \setminus \llbracket \pi \rrbracket_m^0 \\ \llbracket \pi_0 \wedge \pi_1 \rrbracket_m^0 &:= \llbracket \pi_0 \rrbracket_m^0 \cap \llbracket \pi_1 \rrbracket_m^0 \\ \llbracket \pi_0 \vee \pi_1 \rrbracket_m^0 &:= \llbracket \pi_0 \rrbracket_m^0 \cup \llbracket \pi_1 \rrbracket_m^0. \end{aligned}$$

If $s \in \llbracket \pi \rrbracket_m^0$, we will say that π is true or satisfied at s under m , and sometimes write $S, m, s \Vdash^0 \pi$.

Similarly, the *1-step interpretation* $\llbracket \alpha \rrbracket_m^1$ of $\alpha \in 1\text{ML}_\Lambda(A)$ is defined as a subset of TS , with

$$\llbracket \heartsuit_\lambda(\pi_0, \dots, \pi_{n-1}) \rrbracket_m^1 := \lambda_S(\llbracket \pi_0 \rrbracket_m^0, \dots, \llbracket \pi_{n-1} \rrbracket_m^0),$$

and standard clauses applying for $\perp, \wedge, \vee$ and $\neg$. Given a one-step model (S, σ, m) , we write $S, \sigma, m \Vdash^1 \alpha$ for $\sigma \in \llbracket \alpha \rrbracket_m^1$. $\triangleleft$

Remark 7.5 The semantics of the one-step formulas defined above show that, indeed, one-step logic is a way of ‘doing coalgebraic logic without coalgebras’. The link with the interpretation of coalgebraic modal logic is as follows.

Let $\varphi \in \text{ML}_\Lambda$ be some formula of the form $\varphi = \heartsuit_\lambda(\psi_0, \dots, \psi_{n-1})$, and let $\mathbf{Sfor}_0(\varphi) = \{\psi_0, \dots, \psi_{n-1}\}$ denote the set of *direct* subformulas of φ . Take a set A of variables, and some substitution $\tau : A \rightarrow \mathbf{Sfor}_0(\varphi)$ mapping variables to direct subformulas of φ . Now any T -model $\mathbb{S} = (S, \sigma, U)$ induces a A -marking $m = m_{\tau, \mathbb{S}}$ given by

$$m : s \mapsto \{a \in A \mid \mathbb{S}, s \Vdash \tau(a)\}.$$

In other words, we now think of the formulas in $\mathbf{Sfor}_0(\varphi)$ as being atomic. It is then straightforward to verify that we have the following semantic link, for any formula $\alpha \in 1\text{ML}_\Lambda(A)$ such that $\alpha[\tau] = \varphi$:

$$\llbracket \alpha \rrbracket_m^1 = \llbracket \varphi \rrbracket^{\mathbb{S}}. \quad (63)$$

Formulated in terms of one-step models, we have that, for all $s \in S$:

$$S, m, \sigma(s) \Vdash^1 \alpha \text{ iff } \mathbb{S}, s \Vdash \varphi.$$

Notions like one-step satisfiability, validity and equivalence are defined in the obvious way.

Definition 7.6 Let α and α' be one-step formulas. The formula α is *one-step satisfiable* if there is a one-step model (S, σ, m) such that $S, \sigma, m \Vdash^1 \alpha$, and *one-step valid* if $S, \sigma, m \Vdash^1 \alpha$ for all one-step models (S, σ, m) . We say that α' is a *one-step consequence* of α (written $\alpha \models^1 \alpha'$) if $S, \sigma, m \Vdash^1 \alpha$ implies $S, \sigma, m \Vdash^1 \alpha'$, for all one-step models (S, σ, m) , and that α and α' are *one-step equivalent*, notation: $\alpha \equiv^1 \alpha'$, if $\alpha \models^1 \alpha'$ and $\alpha' \models^1 \alpha$. $\triangleleft$

We also need morphisms between one-step frames and models.

Definition 7.7 A *one-step frame morphism* between two one-step frames (S', σ') and (S, σ) is a map $f : S' \rightarrow S$ such that $(Tf)\sigma' = \sigma$. In case such a map satisfies $m' = m \circ f$, for some markings m and m' on S and S' , respectively, viz.,

$$\begin{array}{ccc} S' & \xrightarrow{f} & S \\ & \searrow m' & \swarrow m \\ & PA & \end{array}$$

then we say that f is a *one-step model morphism* from (S', σ', m') to (S, σ, m) . $\triangleleft$

The following proposition, stating that the truth of one-step formulas is invariant under one-step morphisms, is fundamental. We will occasionally refer to this proposition as *naturality*, since this invariance essentially boils down to the naturality of the predicate liftings in Λ .

Proposition 7.8 Let $f : (S', \sigma', m') \rightarrow (S, \sigma, m)$ be a morphism of one-step models over A . Then for every formula $\alpha \in 1ML_\Lambda(A)$ we have

$$S', \sigma', m' \Vdash^1 \alpha \text{ iff } S, \sigma, m \Vdash^1 \alpha.$$

Formulating it differently, for any marking $m : S \rightarrow PA$ and any map $f : S' \rightarrow S$, we have

$$\llbracket \alpha \rrbracket_{m \circ f}^1 = (\check{P}Tf) \llbracket \alpha \rrbracket_m^1. \quad (64)$$

The following proposition states that the meaning of a one-step formula only depends on the variables occurring in it.

Proposition 7.9 Let (S, σ, m) be a one-step model over A , and let $\alpha \in 1ML_\Lambda(A)$ be a one-step formula which belongs to the set $1ML_\Lambda(A')$, for some subset $A' \subseteq A$. Then we have

$$S, \sigma, m \Vdash^1 \alpha \text{ iff } S, \sigma, m^{A'} \Vdash^1 \alpha,$$

where $m^{A'}$ is the A' -marking given by $m^{A'}(s) := m(s) \cap A'$.

Finally, an important role in one-step logic is played by the following rather special one-step models.

Definition 7.10 Given a variable set A , we define the *canonical marking for A* as the A -marking $n_A := \text{id}_{PA} : PA \rightarrow PA$. We say that a one-step model is *A -canonical* if it is of the form (PA, Γ, n_A) , for some $\Gamma \in TPA$. In the sequel we will occasionally denote the marking $\text{id}_{PA} : PA \rightarrow PA$ as n_V . $\triangleleft$

The term ‘canonical’ model is justified by the following proposition, the proof of which is left as an exercise.

Proposition 7.11 *Let $\alpha \in 1\text{ML}_\Lambda(A)$ be a one-step formula. Then α is one-step valid iff α holds at every A -canonical one-step model.*

7.2 Monotone signatures and disjunctive bases

If Λ is a collection of *monotone* predicate liftings, we may further simplify the formulas of one-step modal logic. In this setting we can also introduce the fundamental concepts of disjunctive formulas and disjunctive bases, the latter providing us with a generic framework for obtaining *normal forms* in coalgebraic modal logic.

Definition 7.12 Let $\lambda : \check{P}^n \rightarrow \check{P}T$ be a predicate lifting for the set functor T . We call λ *monotone* if it preserves the subset order in each coordinate, that is, for all sets S and all tuples of pairs of subsets $X_i, Y_i \subseteq S$, we have that

$$\text{if } X_i \subseteq Y_i, \text{ all } i, \text{ then } \lambda_S(X_0, \dots, X_{n-1}) \subseteq \lambda_S(Y_0, \dots, Y_{n-1}).$$

The n -ary predicate lifting λ^∂ defined by putting

$$\lambda_S^\partial(X_0, \dots, X_{n-1}) := TS \setminus \lambda_S(S \setminus X_0, \dots, S \setminus X_{n-1})$$

is called the *boolean dual* of λ . $\triangleleft$

It is straightforward to verify that the boolean dual of a monotone predicate lifting is again monotone.

In the setting where all predicate liftings are monotone, we may generally restrict attention to the positive fragment of one-step modal logic, provided that Λ is closed under taking boolean duals.

Definition 7.13 A *monotone modal signature* for a set functor T is a collection Λ of monotone predicate liftings which is closed under taking boolean duals. $\triangleleft$

Definition 7.14 Fix a monotone modal signature Λ for the set functor T . Given any set E of expressions, we let $\text{PL}^+(E)$ denote the set of *lattice formulas* over E , given by the following grammar:

$$\pi ::= e \in E \mid \perp \mid \top \mid \pi_0 \vee \pi_1 \mid \pi_0 \wedge \pi_1.$$

Where A is a set of propositional variables, we define the set $1\text{ML}_\Lambda^+(A)$ of *positive one-step Λ -formulas* as $1\text{ML}_\Lambda^+(A) := \text{PL}^+(\Lambda(\text{PL}^+(A)))$. $\triangleleft$

A more direct characterisation of the language $1\text{ML}_\Lambda(A)$ is that it consists of those negation-free $\text{ML}_\Lambda(A)$ -formulas in which every variable $a \in A$ occurs in the scope of exactly one modality.

For positive one-step formulas we have the following monotonicity property.

Proposition 7.15 *Assume that Λ is a monotone modal signature for a set functor T , let (S, σ) be a one-step frame, and let $m, m' : S \rightarrow PA$ be two markings such that $m(s) \subseteq m'(s)$, for all $s \in S$. Then for any formula $\alpha \in 1\text{ML}_\Lambda^+(A)$ we have that*

$$S, \sigma, m \Vdash^1 \alpha \text{ implies } S, \sigma, m' \Vdash^1 \alpha.$$

We now introduce the fundamental concept of a disjunctive basis. First we consider the notion of a disjunctive formula.

Definition 7.16 A one-step formula $\alpha \in 1\text{ML}_\Lambda^+(A)$ is called *disjunctive* if for every one-step model (S, σ, m) such that $S, \sigma, m \Vdash^1 \alpha$ there is a one-step frame (S', σ') together with a one-step frame morphism $h : (S', \sigma') \rightarrow (S, \sigma)$ and a marking $m' : S' \rightarrow PA$, such that:

1. $S', \sigma', m' \Vdash^1 \alpha$;
2. $m'(s') \subseteq m(h(s'))$, for all $s' \in S'$;
3. $|m'(s')| \leq 1$, for all $s' \in S'$.

We sometimes refer to the one-step frame (S', σ') together with the map h as a *cover* of (S, σ) , and to the one-step model (S', σ', m') together with the map h as a *separating cover* of (S, σ, m) for α . $\triangleleft$

The intuition behind disjunctive formulas is that, in a certain sense, they never ‘force’ two distinct propositional variables to be true together, that is, any one-step model for a disjunctive formula δ in $1\text{ML}_\Lambda^+(A)$ is covered by one in which every point satisfies at most one propositional variable from A .

Example 7.17 (1) A trivial example of a disjunctive formula is $\bigcirc a$ for $a \in A$, where we recall that $\bigcirc$ was the next-time modality viewed as a predicate lifting for the identity functor Id . A one-step model for this functor is a triple S, s, m consisting of a set S , an element $s \in S$ and a marking $m : S \rightarrow PA$. Then $S, s, m \Vdash^1 \bigcirc a$ if, and only if, $a \in m(s)$. But then, no elements in S besides s are relevant to the evaluation of $\bigcirc a$, and for s we can just forget about all other variables: set $m'(s) = \{a\}$ and $m'(v) = \emptyset$ for all $s \in S \setminus \{v\}$. We have $S, s, m' \Vdash^1 \bigcirc a$, $m'(v) \subseteq m(v)$ and $|m'(v)| \leq 1$ for all $v \in S$.

(2) For an example of a one-step formula that is *not* disjunctive, consider $\Box a \wedge \Diamond b$ (where $\Diamond$ and $\Box$ are the standard modalities for Kripke structures, that is, coalgebras for the power set functor P). Observe that a one-step model for this functor is a triple (S, σ, m) with $\sigma \subseteq S$. It should be obvious that for the formula $\Box a \wedge \Diamond b$ to hold at such a structure, σ needs to have an element s where b holds, while at the same time *every* element of σ , including s , must satisfy a . There is no escape here: we can only have $(S, \sigma, m) \Vdash^1 \Box a \wedge \Diamond b$ if there is an element s making *both* a and b true.

(3) This is very different if we consider the typical disjunctive formulas for basic modal logic, which are of the form $\nabla\{a_1, \dots, a_n\}$, or, in the language with $\Box$ and $\Diamond$:

$$\Diamond a_1 \wedge \dots \wedge \Diamond a_n \wedge \Box(a_1 \vee \dots \vee a_n)$$

For example, consider the formula $\nabla\{a, b, c\}$, which is true in a one-step model (S, s, m) with $S = \{u, v\}$, $m(u) = \{a, b\}$ and $m(v) = \{b, c\}$. Clearly there is no way to simply shrink the marking m to a separating marking m' so that $(S, s, m) \Vdash^1 \nabla\{a, b, c\}$: there are too many pigeons and too few pigeon holes — so the obvious solution is to make more pigeon holes! One way to do this is to “split” the points in $\{u, v\}$ so that we make room for each variable to be witnessed at a separate point. More formally, let $S' = \{(u, a), (u, b), (v, b), (v, c)\}$ and define $m' : S' \rightarrow \{a, b, c\}$ via the projection on $\{a, b, c\}$, i.e., $m'(u, a) = \{a\}$, etc. This one-step model satisfies $\nabla\{a, b, c\}$ and it has the obvious covering map h being the projection π_S on S .

(4) In fact, all nabla formulas are disjunctive. That is, let T be a standard and smooth set functor; given a set A and an element $\alpha \in T_\omega A$ we may think of $\nabla\alpha$ as a one-step formula. It is not hard to show that every one-step formula of this form is disjunctive.

As we will see later on, disjunctive formulas have very attractive properties, and so modal signatures with a rich supply of disjunctive formulas are of interest. Example 7.17(3) shows that not *every* one-step formula is disjunctive. It does make sense to consider modal signatures with a ‘critical’ amount of disjunctive formulas in the form of a *disjunctive basis*. Roughly speaking, the existence of a *disjunctive basis* means that the coalgebraic modal logic satisfies some version of the *modal distributive laws* that we proved for the ∇ -logic in Chapter 5. For more details, we need some definitions.

Definition 7.18 We will consider, next to a (finite) collection A of variables, every element $B \in P_\omega A$ as a (formal) variable $\underline{B}$, and we shall be interested in the substitution

$$\bigwedge_A : P_\omega A \rightarrow A$$

that replaces each formal variable $\underline{B}$ with its corresponding conjunction $\bigwedge B \in \text{PL}^+(A)$. $\triangleleft$

Intuitively then, Λ having a disjunctive basis means that there are sufficiently many disjunctive formulas so that for every formula $\alpha \in \text{1ML}_\Lambda(A)$ there is a disjunctive formula $\delta \in \text{1ML}_\Lambda(P_\omega A)$ such that

$$\alpha \equiv \delta[\bigwedge_A]. \quad (65)$$

Observe that if we take α to be a conjunction of disjunctive formulas (65) expresses a law stating that conjunctions distribute over disjunctive formulas.

Example 7.19 $\blacktriangleright$ example 1: next time operator

$\blacktriangleright$ example 2: $\Box a \wedge \Diamond b$

We formulate the definition of a disjunctive basis in such a way that we only need to check the existence of disjunctive formulas satisfying (65) in a small but representative number of cases.

Definition 7.20 Let D be an assignment of a set of positive one-step formulas $D(A) \subseteq 1ML_{\Lambda}^+(A)$ for all sets of variables A . Then D is called a *disjunctive basis* for Λ if each formula in $D(A)$ is disjunctive, and the following conditions hold:

1. $D(A)$ contains $\top$ and is closed under finite disjunctions (in particular, also $\perp = \bigvee \emptyset \in D(A)$).
2. D is *distributive over Λ* : for every one-step formula of the form $\heartsuit_{\lambda}\bar{\pi}$ in $1ML_{\Lambda}^+(A)$ there is a formula $\delta \in D(P(A))$ such that $\heartsuit_{\lambda}\bar{\pi} \equiv \delta[\bigwedge_A]$.
3. D *admits a binary distributive law*: for any two formulas $\alpha \in D(A)$ and $\alpha' \in D(A')$, there is a formula $\gamma \in D(P^{*2}(A \cup A'))$ such that $\alpha \wedge \alpha' \equiv \gamma[\bigwedge_{A \cup A'}]$.

Here $P^{*2}(A \cup A')$ is the collection of subsets C of $A \cup A'$ such that either $|C| = 1$ or $|C| = 2$ and C overlaps with both A and A' . $\triangleleft$

The key property of disjunctive bases is captured by the following normal form theorem, which is easy to derive from the definition.

Proposition 7.21 *Suppose D is a disjunctive basis for Λ . Then for every one-step formula $\alpha \in 1ML_{\Lambda}^+(A)$ there is a formula $\delta \in D(P(A))$ such that (65) holds: $\alpha \equiv \delta[\bigwedge_A]$.*

► examples to be added

7.3 One-step derivation systems: sequent style

In this section and the next we will develop some natural *logic* at the level of one-step formulas; that is, we will introduce one-step derivation systems and discuss the notions of one-step soundness and completeness pertaining to such systems. As we will see in more detail in the next Chapter, the basic idea is that coalgebraic derivation systems consist of so-called *one-step rules*, on top of a classical propositional calculus.

The derivation system in this section will be based on *sequents*, that is, finite sets of formulas. We read sequents *disjunctively*, that is, we think of a sequent Γ as representing the formula $\bigvee \Gamma$, and the idea of a derivation of the sequent Γ is to provide a formal proof that this disjunction is valid. Following standard conventions in proof theory we will often denote a singleton sequent by its unique member (that is, write φ instead of $\{\varphi\}$), and use the comma for the union symbol (that is, let Γ, Δ denote the union $\Gamma \cup \Delta$ of the two sequents Γ and Δ).

Remark 7.22 Our set-up, which is sometimes referred to as Tait-style, differs from the slightly more common approach in which sequents are *pairs* of sets (or even multi-sets) of formulas, with the pair (Γ, Δ) , usually written as $\Gamma \vdash \Delta$ or $\Gamma \Rightarrow \Delta$ representing the formula $\bigwedge \Gamma \rightarrow \bigvee \Delta$. The point is that, given our classical (boolean) context we can represent the two-sided sequent $\Gamma \Rightarrow \Delta$ as $\emptyset \Rightarrow \bar{\Gamma}, \Delta$, where $\bar{\Gamma} := \{\bar{\gamma} \mid \gamma \in \Gamma\}$ is the set of negations of the formulas in Γ . We may then subsequently drop the (fixed) antecedent $\emptyset$ and focus on the Tait-style sequent $\bar{\Gamma}, \Delta$. Conversely, a Tait-style sequent Σ represents the consequent (right hand side) of the two-sided sequent $\emptyset \Rightarrow \Sigma$.

Based on these correspondences one may develop the proof theory of two-sided sequent system along the same lines as what we doing here for one-sided sequent systems.

Remark 7.23 As a matter of fact, we could also have pursued a *tableau-style* approach where the idea is to focus on the antecedent part (left hand side) of a two-sided sequent. Such sequents are then to be read conjunctively, and the idea of a derivation is to *refute* the sequent, i.e., to provide a formal proof that the sequent is not satisfiable.

7.3.1 Rules and derivations

We shall formulate derivation rules in terms of propositional and rank-1 formulas over A . Accordingly, we shall refer to finite sets of propositional formulas, (rank-1 formulas, respectively) as *propositional* and *rank-1* sequents.

The key concept in coalgebraic derivation system is that of a *one-step derivation rule*.

Definition 7.24 A *derivation rule* is a pair $d = (\text{Prem}(d), \text{conc}(d))$, where $\text{Prem}(d)$ is a set of sequents, the *premises* of d , and $\text{conc}(d)$ is a sequent, the *conclusion* of d . If $\text{Prem}(d) = \emptyset$ we say that d is *axiomatic* and we call $\text{conc}(d)$ an *axiom*. A *derivation system* or *calculus* is nothing but a set of derivation rules.

A derivation rule is *propositional* if all of its premises and its conclusion are propositional sequents, whereas a *one-step* rule is given by a set of propositional premises and a rank-1 conclusion. In particular, one-step axioms are sets of one-step formulas. $\triangleleft$

A derivation rule d with $\text{Prem}(d) = \{\Delta_1, \dots, \Delta_{n-1}\}$ and $\text{conc}(d) = \Delta_0$ is often represented as $\Delta_1, \dots, \Delta_{n-1} / \Delta_0$ or as

$$\frac{\Delta_1 \quad \cdots \quad \Delta_{n-1}}{\Delta_0} d$$

Example 7.25 Examples of propositional rules are Modus Ponens $(\{a, a \rightarrow b\}, b)$ and Ex Falso Quodlibet $(\{\perp\}, a)$.

In ordinary modal logic (that is, where $T = P$ and we take both the diamond and the box modality as primitives), we can present the usual modal rule as follows:

$$\frac{a, b_0, \dots, b_{n-1}}{\Box a, \Diamond b_0, \dots, \Diamond b_{n-1}} M$$

Other examples of one-step rules are the congruence rule C_λ :

$$\frac{a_0 \leftrightarrow b_0 \quad \cdots \quad a_{n-1} \leftrightarrow b_{n-1}}{\heartsuit_\lambda(a_0, \dots, a_{n-1}) \leftrightarrow \heartsuit_\lambda(b_0, \dots, b_{n-1})} C_\lambda$$

and the monotonicity rule M_λ

$$\frac{a_0 \rightarrow b_0 \quad \cdots \quad a_{n-1} \rightarrow b_{n-1}}{\heartsuit_\lambda(a_0, \dots, a_{n-1}) \rightarrow \heartsuit_\lambda(b_0, \dots, b_{n-1})} M_\lambda$$

that we will associate with an n -ary modality $\heartsuit_\lambda$.

Remark 7.26 As mentioned already, the basic idea in the development of proof theory for coalgebraic modal logic is to consider derivation systems that consist of a *fixed* set of propositional derivation rules, together with a *specific*, purpose-built collection of one-step rules. For instance, in the setting of one-sided sequent systems, we will consider the following propositional rules for the connectives $\top$, $\neg$ and $\wedge$:

$$\frac{}{a, \neg a} \text{Ax0} \quad \frac{}{\top} \text{Ax1} \quad \frac{a}{\neg \neg a} \text{R}_{\neg} \quad \frac{a \quad b}{a \wedge b} \text{R}_{\wedge} \quad \frac{\neg a, \neg b}{\neg(a \wedge b)} \text{R}_{\neg \wedge}$$

Given a set $\mathbf{G}$ of one-step rules, we obtain the system $\mathbf{G}^+$ by adding a suitable set of propositional rules, including the ones given above and similar ones for the other primitive propositional connectives.

To introduce *derivations*, we need to say how we can use substitutions to define instances of derivation rules.

Definition 7.27 A *substitution* is a map $\tau : A \rightarrow \text{ML}_{\Lambda}(A)$. We will use the notation φ/a for the substitution that maps the variable a to the formula φ (and is the identity on the set of remaining variables). A substitution τ is *propositional* or *rank-0* if $\tau : A \rightarrow \text{PL}(A)$, and *rank-1* if $\tau : A \rightarrow 1\text{ML}_{\Lambda}(A)$. A *renaming* is a simple substitution $\tau : A \rightarrow A$ that replaces variables with variables.

A substitution τ naturally induces a translation $[\tau]$ mapping $\text{ML}_{\Lambda}(A)$ -formulas to $\text{ML}_{\Lambda}(A)$ -formulas. For this translation we shall use postfix notation, $\varphi[\tau]$ denoting the result of applying the substitution τ to the formula φ . For a sequent Φ we will write $\Phi[\tau] := \{\varphi[\tau] \mid \varphi \in \Phi\}$.

<

Remark 7.28 It is easy to see that if τ is a propositional substitution, we obtain $[\tau] : \text{PL}(A) \rightarrow \text{PL}(A)$ and $[\tau] : 1\text{ML}_{\Lambda}(A) \rightarrow 1\text{ML}_{\Lambda}(A)$. On the other hand, if τ is rank-1, then $[\tau]$ maps rank-0 formulas to rank-1 formulas, $[\tau] : \text{PL}(A) \rightarrow 1\text{ML}_{\Lambda}(A)$. In the sequel we will use these observations without explicit notice.

Roughly speaking, an *instance* of a derivation rule is a version of that rule where we uniformly replace every formula in the rule with some substitution instance, and we (possibly) add context formulas.

Definition 7.29 To define the *instances* of a derivation rule $\mathbf{d} = \Delta_1, \dots, \Delta_{n-1} / \Delta_0$ is, we make a case distinction:

- If $\mathbf{d}$ is propositional, then any instance of $\mathbf{d}$ can be represented as a rule

$$\frac{\Gamma, \Delta_1[\tau] \quad \dots \quad \Gamma, \Delta_{n-1}[\tau]}{\Gamma, \Delta_0[\tau]} \mathbf{d}$$

where τ is a substitution and Γ is a set of formulas.

- If d is a one-step rule, then any instance of d can be represented as a rule

$$\frac{\Delta_1[\tau] \cdots \Delta_{n-1}[\tau]}{\Gamma, \Delta_0[\tau]} d$$

where τ is a rank-0 substitution.

In both cases we refer to the formulas in $\Delta_0[\tau]$ as the *principal* formulas of the rule instance, and the ones in Γ as the *context* formulas. $\triangleleft$

Note the distinction between instances of propositional and one-step rules when it comes to the role of the context formulas: in the case of propositional rules, the context formulas in the conclusion are also part of each premise, whereas in the case of a one-step rule the context only appears in the conclusion.

Definition 7.30 Let $\mathbf{D}$ be a derivation system. A *derivation* in $\mathbf{D}$ is a structure (W, C, L, r) , where (W, C, r) is a well-founded tree¹⁰ with root r , and L is a map *labelling* every node w of W with a sequent¹¹ $L(w)$, in such a way that for every inner node $t \in W$, the pair $(\{L(u) \mid u \in C(t)\}, L(t))$ is an instance of a derivation rule in $\mathbf{D}$.

Given an $\mathbf{D}$ -derivation $\mathcal{D} = (W, C, L, r)$, we call a sequent Σ an *assumption* of $\mathcal{D}$ if $\Sigma = L(t)$ for some *leaf* t , but Σ is not an instance of an axiom; we denote the set of assumptions of $\mathcal{D}$ by $\text{Ass}(\mathcal{D})$. We refer to the sequent $L(r)$ labelling the root of the tree as the *result* of $\mathcal{D}$.

A sequent $\Sigma \in 1\text{ML}_\Lambda(A)$ is *derivable* from a collection $\mathcal{A}$ of sequents in a derivation system $\mathbf{D}$, notation: $\mathcal{A} \vdash_{\mathbf{D}} \Sigma$ if Σ is the result of an $\mathbf{D}$ -derivation $\mathcal{D}$ with $\text{Ass}(\mathcal{D}) \subseteq \mathcal{A}$. If $\emptyset \vdash_{\mathbf{D}} \Sigma$ we simply write $\vdash_{\mathbf{D}} \Sigma$ and we say that Σ is *derivable*. $\triangleleft$

Remark 7.31 Derivations of one-step sequents in one-step derivation systems have a rather specific shape. The tree of such a derivation can be partitioned into an ‘upper’ and a ‘lower’ part, where the nodes of the upper (lower) part are all labelled with rank-0 (respectively, rank-1) sequents. Formulated differently, on every branch $t_0 t_1 \cdots t_n$ from the root $r = t_0$ to a leaf $l = t_n$ of the tree, we either have $L(t_i) \subseteq 1\text{ML}_\Lambda(A)$, for every i , or else there is a (necessarily unique) index $k \geq 0$ such that $L(t_i) \subseteq 1\text{ML}_\Lambda(A)$ for all i with $0 \leq i \leq k$, while $L(t_i) \subseteq \text{PL}(A)$ for all i with $k < i \leq n$. In the latter case $L(t_k)$ is the only sequent on the mentioned branch that is obtained by the application of a *one-step* rule — all other sequents are the conclusion of an application of a *propositional* rule.

¹⁰Trees and related notions are defined in Definition A.8.

¹¹Sometimes we will also need to make explicit the actual rule that has been applied at the node.

7.3.2 One-step soundness and completeness

We now introduce the two key concepts in the one-step approach towards derivation systems for coalgebraic modal logics: *one-step soundness* and *one-step completeness*.

We may restrict attention to one-step rules in a simplified form. Recall that $\text{Lit}(A)$ is the set of *literals* over A , that is, variables and negations of variables in A .

Definition 7.32 A *modal rule* is a one-step rule d such that every $\Delta \in \text{Prem}(d)$ satisfies $\Delta \subseteq \text{Lit}(A)$ while $\text{conc}(d) \subseteq \text{Lit}(\Lambda(A))$. Furthermore we require that every variable that occurs in one the premises also appears in the conclusion. $\triangleleft$

Remark 7.33 The restriction to modal rules can be justified by the observation that every one-step derivation system $\mathbf{G}$ can be replaced with a modal derivation system $\mathbf{M}_{\mathbf{G}}$ that is equivalent to $\mathbf{G}$ in the sense that the same set of one-step rules are derivable in the two systems (extended with propositional rules). More details of this equivalence will be provided later.

The definition of one-step soundness is not surprising.

Definition 7.34 A modal rule d is *one-step sound* if, for every marking $m : X \rightarrow PA$ such that $\llbracket \Delta \rrbracket_m^0 = X$ for each $\Delta \in \text{Prem}(d)$ we have $\llbracket \text{conc}(d) \rrbracket_m^1 = TX$. A set $\mathbf{G}$ is *one-step sound* if every rule in $\mathbf{G}$ is so. $\triangleleft$

Remark 7.35 We leave it for the reader that all propositional-substitution instances of sound one-step rules are sound.

The definition of one-step completeness is somewhat more involved.

Definition 7.36 A set $\mathbf{G}$ of modal rules is *one-step complete* if, for every $\Gamma \subseteq \text{Lit}(\Lambda(A))$ and every marking $m : X \rightarrow PA$ such that $\llbracket \Gamma \rrbracket_m^1 = TX$, there is some rule $d = \Gamma_1, \dots, \Gamma_n / \Gamma_0$ and some renaming $\tau : A \rightarrow A$ such that $\Gamma_0[\tau] \subseteq \Gamma$ and $\llbracket \Gamma_i[\tau] \rrbracket_m^0 = X$, for all i . $\triangleleft$

Remark 7.37 The above definition is not the only way to make the concept of one-step completeness precise. Here are two alternative (but basically equivalent) approaches.

First of all, we could get rid of renamings in the definition. That is, we could say that a set $\mathbf{G}$ of modal rules is one-step complete if, for every $\Gamma \subseteq \text{Lit}(\Lambda(A))$ and every marking $m : X \rightarrow PA$ such that $\llbracket \Gamma \rrbracket_m^1 = TX$, there is some rule $d = \Gamma_1, \dots, \Gamma_n / \Gamma_0$ such that $\Gamma_0 \subseteq \Gamma$ and $\llbracket \Gamma_i \rrbracket_m^0 = X$, for all i . The disadvantage of this definition is that for any rule in the system we would also need to have each of its renamings.

Moving in the opposite direction, we could call a set $\mathbf{G}$ of modal rules one-step complete if, for every $\Gamma \subseteq \text{Lit}(\Lambda(A))$ and every marking $m : X \rightarrow PA$ such that $\llbracket \Gamma \rrbracket_m^1 = TX$, there are rank-0 sequents $\Gamma_1, \dots, \Gamma_n$ such that Γ is *derivable* from $\Gamma_1, \dots, \Gamma_n$ using the rules in $\mathbf{G}$, in combination with propositional logic. This definition would lead to smaller derivation systems, but a drawback is that it involves propositional reasoning and thus somewhat obscures the fact that the notion of one-step completeness is really about one-step rules.

Example 7.38 We look at standard modal logic (that is, $T = P$), taking the diamond operator $\Diamond$ as the single modality. We claim that the set of rules of the form

$$\frac{\neg a_0, a_1, \dots, a_n}{\neg \Diamond a_0, \Diamond a_1, \dots, \Diamond a_n} \mathbf{m}_n$$

is one-step complete.

To see this, assume that $\llbracket \Gamma \rrbracket_m^1 = TX$, for some marking $m : X \rightarrow PA$ and some sequent $\Gamma \subseteq \text{Lit}(\{\Diamond a \mid a \in A\})$, say, $\Gamma = \{\neg \Diamond b_0, \dots, \neg \Diamond b_{k-1}, \Diamond c_0, \dots, \Diamond c_{n-1}\}$.

We claim that there is an i for which $\llbracket \neg b_i, c_0, \dots, c_{n-1} \rrbracket_m^0$ is the full set X . To arrive at a contradiction, assume that $\llbracket \neg b_i, c_0, \dots, c_{n-1} \rrbracket_m^0 \neq X$, for every i . Then for each i we may take some object $x_i \notin \llbracket \neg b_i, c_0, \dots, c_{n-1} \rrbracket_m^0$. Defining $Y := \{x_i \mid 0 \leq i < k\}$, we find $Y \notin \llbracket \Gamma \rrbracket_m^1$, contradicting our earlier assumption. This means that our claim holds indeed.

Let Δ_0 and Δ_1 be the sequents $\Delta_0 := \{\neg a_0, a_1, \dots, a_n\}$ and $\Delta_1 := \{\neg \Diamond a_0, \Diamond a_1, \dots, \Diamond a_n\}$, so that Δ_0/Δ_1 is the rule $\mathbf{m}_n$. Furthermore, consider the renaming τ given by $\tau : a_0 \mapsto b_i$ and $\tau : a_j \mapsto c_{j-1}$ for $j \geq 1$. We then have $\llbracket \Delta_0[\tau] \rrbracket = X$ and $\Delta_1[\tau] \subseteq \Delta$, as required by the definition of one-step completeness.

The following Theorem states that we can *always* find a one-step complete set of one-step rules — for every set functor and every modal signature for that functor.

Theorem 7.39 *Let Λ be a modal signature for a set functor T . Then the collection S_Λ of all sound one-step rules is one-step complete.*

7.4 One-step derivation systems: Hilbert-style

In this section we continue our exploration of one-step derivation systems, but now we concentrate on so-called *Hilbert-style systems*. Characteristically, these manipulate single formulas, rather than sequents.

Definition 7.40 A *Hilbert calculus* is a derivation system consisting of one-step rules in which the premises and conclusion of every rule are all singletons. A *(Hilbert) axiomatization* is a Hilbert calculus consisting of axiomatic rules only. $\triangleleft$

A Hilbert-style derivation rule will usually be presented as a pair $R = (\Pi_R, \gamma_R)$ where Π_R is a set of formulas and γ_R is a formula – that is, we replace the singleton sequents with their unique formula members. Also, we will identify a Hilbert axiomatisation $\mathbf{H}$ with the set $\{\gamma_R \mid R \in \mathbf{H}\}$ of its axioms.

Remark 7.41 The restriction to singleton sequences (and thus, equivalently, to formulas) extends to derivations well. As a consequence of this, in Hilbert-style derivations we will only consider rule instances where the context is *empty*.

Convention 7.42 Throughout this text we will assume an arbitrary but fixed derivation system $\mathbf{C}$ (consisting of propositional axioms and rules) which is sound and complete for classical propositional logic.

Definition 7.43 Given a derivation system $\mathbf{H}$, we let $\mathbf{H}^+$ denote the extension of $\mathbf{H}$ with (1) all axioms and rules from $\mathbf{C}$ and (2) the congruence rule (C_λ) , for every $\lambda \in \Lambda$. $\triangleleft$

In Hilbert calculi the tendency is to prefer axioms to non-axiomatic rules. In fact one often fixes a minimal set of derivation rules and focuses on one-step axiomatizations on top of these. This approach is justified by Theorem 7.44 below which states that in fact we may completely restrict attention to derivation systems in which the congruence rule is the *only* non-axiomatic one-step rule.

Theorem 7.44 Let Λ be some set of predicate liftings for a set functor T , and let $\mathbf{H}$ be a one-step Hilbert-style derivation system. Then there is an axiomatization $\mathbf{A}$ such that $\mathbf{H}^+$ and $\mathbf{A}^+$ derive the same one-step formulas.

Proof. Clearly it suffices to show that we can replace every one-step rule R with an axiom α_R which is equivalent to R in the sense that R and α_R are derivable from one another by propositional reasoning and the congruence rule. Without loss of generality we may assume the premises of R to be satisfiable; if not, we may just reply R with the axiom $\top$.

Our proof is then based on an interesting property of classical propositional logic: The claim below states that every satisfiable propositional formula has a so-called *projective unifier*.

Claim 1 Let $\pi \in \text{PL}(A)$ be satisfiable. Then π has a projective unifier, that is, a substitution $\tau : A \rightarrow \text{PL}(A)$ such that (1) $\pi \models \tau$ and (2) $\pi \models a \leftrightarrow \tau(a)$, for all $a \in A$.

PROOF OF CLAIM Given a propositional formula ρ and a valuation $U : A \rightarrow 2$ we write $\llbracket \rho \rrbracket_U \in \{0, 1\}$ to denote the truth value of ρ under U .

Since π is satisfiable, there is a valuation $V : A \rightarrow 2$ with $\llbracket \pi \rrbracket_V = 1$. Let τ be given by

$$\tau(a) := \begin{cases} \pi \wedge a & \text{if } V(a) = 1 \\ \pi \rightarrow a & \text{if } V(a) = 0. \end{cases}$$

It is straightforward to verify that $\pi \models a \leftrightarrow \tau(a)$, for all $a \in A$. To check that $\pi[\tau]$ is a tautology, consider an arbitrary valuation $U : A \rightarrow 2$. We claim that $\llbracket \pi[\tau] \rrbracket_U = 1$.

To see this, we make a case distinction. If $\llbracket \pi \rrbracket_U = 1$, we find by part (2) that $\llbracket a \rrbracket_U = \llbracket \tau(a) \rrbracket_U$, and from this an easy induction shows that $\llbracket \rho \rrbracket_U = \llbracket \rho[\tau] \rrbracket_U$, for any formula $\rho \in \text{PL}(A)$. In particular we have $\llbracket \pi \rrbracket_U = \llbracket \pi[\tau] \rrbracket_U$ and since $\llbracket \pi \rrbracket_U = 1$ this gives $\llbracket \pi[\tau] \rrbracket_U = 1$, as required. On the other hand, if $\llbracket \pi \rrbracket_U = 0$ then by definition of τ , no matter whether $V(a) = 0$ or $V(a) = 1$, we always have $\llbracket a \rrbracket_U = \llbracket \tau(a) \rrbracket_U$. From this again an easy induction shows that $\llbracket \rho \rrbracket_U = \llbracket \rho[\tau] \rrbracket_U$, for any formula $\rho \in \text{PL}(A)$. In particular then for $\rho = \pi$ we find that $\llbracket \rho[\tau] \rrbracket_U = \llbracket \rho \rrbracket_U = 0$. $\triangleleft$

Let R be the rule Π/γ , and let τ be the substitution we obtain by applying the claim to the formula $\bigwedge \Pi$. Then we define

$$\alpha_R := \gamma[\tau].$$

and we claim that R and α_R are mutually derivable.

To see this, by the first property of the claim the formula $\pi[\tau]$ is derivable, for each premise $\pi \in \Pi$. From this we can derive α_R by one application of R . Conversely, take the set Π as assumptions, so that we can derive $\bigwedge \Pi$ by propositional reasoning. Then by the second property of the claim we may derive $a \leftrightarrow \tau(a)$, for all propositions a . But then we may easily derive $\alpha \leftrightarrow \alpha[\tau]$ by applications of the congruence rule, and since $\alpha[\tau]$ is now an axiom we obtain α . QED

Example 7.45 Consider the monotonicity rule $a \rightarrow b / \heartsuit a \rightarrow \heartsuit b$, and observe that there are three valuations making the premise $a \rightarrow b$ true. In the table below we list these, together with the induced substitutions.

V	a	b	$\tau(a)$	$\tau(b)$
V_0	0	0	$(a \rightarrow b) \wedge a \equiv a \wedge b$	$(a \wedge b) \wedge b \equiv b$
V_1	0	1	$(a \rightarrow b) \wedge a \equiv a \wedge b$	$(a \wedge b) \rightarrow b \equiv a \vee b$
V_2	1	1	$(a \rightarrow b) \rightarrow a \equiv a$	$(a \wedge b) \rightarrow b \equiv a \vee b$

Applying these three substitutions to the conclusion $\heartsuit a \rightarrow \heartsuit b$ we obtain three axioms which may replace the monotonicity rule: $\heartsuit(a \wedge b) \rightarrow \heartsuit b$, $\heartsuit(a \wedge b) \rightarrow \heartsuit(a \vee b)$ and $\heartsuit a \rightarrow \heartsuit(a \vee b)$.

We now move to the concepts of one-step soundness and completeness for Hilbert-style systems. The definition of soundness is the same as before.

Definition 7.46 Let Λ be a set of predicate liftings for a set functor T . A one-step rule $R = \Pi/\gamma$ for Λ is *one-step sound* if

$$\bigcap_{\pi \in \Pi} \llbracket \pi \rrbracket_m^0 = S \text{ implies } \llbracket \gamma \rrbracket_m^1 = TS,$$

for all sets S and all markings $m : S \rightarrow PA$. A derivation system $\mathbf{H}$ is *one-step sound* if all of its derivation rules are one-step sound. $\triangleleft$

The definition of one-step completeness is slightly more involved than in the case for sequent systems.

Definition 7.47 Let Λ be a set of predicate liftings for a set functor T . We say that $\pi \in \text{PL}(A)$ is a *true (propositional) fact* of a marking $m : S \rightarrow PA$ if $\llbracket \pi \rrbracket_m^0 = S$; we let $\text{TPF}(m)$ denote the collection of all these facts.

A one-step derivation system $\mathbf{H}$ for Λ is *one-step complete* if for every marking $m : S \rightarrow PA$, and every $\alpha \in \text{1ML}_\Lambda(A)$ we have that

$$\llbracket \alpha \rrbracket_m^1 = TS \text{ implies } \text{TPF}(m) \vdash_{\mathbf{H}^+}^1 \alpha,$$

i.e., all formulas that are one-step true with respect to m are derivable from the true propositional facts of m in the *extended* derivation system $\mathbf{H}^+$. $\triangleleft$

Remark 7.48 Note that in the case of Hilbert-style derivation systems our definition of one-step completeness uses the approach sketched in Remark 7.37, rather than the one of Definition 7.36.

Example 7.49 As an example of a one-step complete derivation system, consider the modal signature $\{\Box\}$ for the powerset functor P (i.e., we look at standard modal logic where we take the $\Box$ modality as primitive). In the sequel we will continue writing T instead of P , however, in order to clarify the role of the functor in the argument. We claim that the axiomatisation

$$\mathbf{K} := \{\Box \top, \Box(a \wedge b) \leftrightarrow \Box a \wedge \Box b\}$$

is one-step sound and complete. Leaving soundness as an exercise for the reader, we prove one-step completeness here.

Let m be a A -marking on some set S , and let α be a one-step formula such that $\llbracket \alpha \rrbracket_m^1 = TS$. We will show that α is derivable from the true propositional facts of m :

$$\text{TPF}(m) \vdash_{\mathbf{K}}^1 \alpha.$$

Since we have the full power of classical propositional logic at our disposal, we may without loss of generality assume that α is in conjunctive normal form, i.e., $\alpha = \bigwedge_{\beta \in B} \beta$ for some finite set B , where each β is of the form

$$\beta = \bigvee_{i \in I} \Box \pi_i \vee \bigvee_{j \in J} \neg \Box \tau_j$$

for some finite index sets I and J , and where all $\pi_i, \tau_j \in \text{PL}(A)$. Clearly it suffices to prove that we can derive each conjunct of α from the true propositional facts of m , so fix such a conjunct $\beta \in B$. Obviously it follows from $\llbracket \alpha \rrbracket_m^1 = TS$ that $\llbracket \beta \rrbracket_m^1 = TS$.

We first claim that

$$I \neq \emptyset. \tag{66}$$

To see this, assume for contradiction that $I = \emptyset$, so that $\beta = \bigvee_{j \in J} \neg \Box \tau_j$. Consider the empty set $\emptyset \in TS$, and observe that $S, m, \emptyset \Vdash^1 \Box \pi$ for *all* rank-0 formulas $\pi \in \text{PL}(A)$; so in particular, $S, m, \emptyset \Vdash^1 \Box \tau_j$ for all $j \in J$. But this means $S, m, \emptyset \not\Vdash^1 \beta$, which clearly contradicts the earlier observation that $\llbracket \beta \rrbracket_m^1 = TS$. This finishes the proof of (66).

Our second claim is that

$$\text{at least one of the formulas } \xi_i \text{ is a true propositional fact of } m, \quad (67)$$

where we define

$$\xi_i = \pi_i \vee \bigvee_{j \in J} \neg \tau_j.$$

To prove (67), assume for contradiction that none of formulas ξ_i belongs to the set $TPF(m)$. Then there are states $(s_i)_{i \in I}$ such that $s_i \notin \llbracket \xi_i \rrbracket_m^0$ for each $i \in I$. So $S, m, s_i \not\Vdash^0 \pi_i$, while $S, m, s_i \Vdash^0 \tau_j$ for all j . Now consider the set $\sigma := \{s_i \mid i \in I\} \in PS$. It is immediate by the semantics of $\Box$ that $S, m, \sigma \not\Vdash^1 \Box \pi_i$ for any $i \in I$, while at the same time $S, m, \sigma \Vdash^1 \Box \tau_j$ for all $j \in J$. Clearly then we find that $S, m, \sigma \not\Vdash^1 \beta$, which provides the desired contradiction with our assumption that $S, m, \sigma \Vdash^1 \alpha$. This proves (67).

To finish the one-step completeness proof, assume $\xi_i = \pi_i \vee \bigvee_{j \in J} \neg \tau_j \in TPF(m)$, then so is the (propositionally equivalent) formula

$$\bigwedge_{j \in J} \tau_j \rightarrow \pi_i.$$

Our third claim is that from the above formula we can derive its ‘boxed version’

$$\bigwedge_{j \in J} \Box \tau_j \rightarrow \Box \pi_i,$$

by some propositional reasoning, applications of the congruence rule and of the axioms $\Box(a \wedge b) \leftrightarrow \Box a \wedge \Box b$ and $\Box \top$, followed by further propositional reasoning. We leave the details for the reader.

Finally, we can use propositional reasoning to show that the formula $\bigwedge_{j \in J} \Box \tau_j \rightarrow \Box \pi_i$ is equivalent to

$$\Box \pi_i \vee \bigvee_{j \in J} \neg \Box \tau_j,$$

and therefore implies

$$\beta = \bigvee_{i \in I} \Box \pi_i \vee \bigvee_{j \in J} \neg \Box \tau_j.$$

In other words, we have established that

$$TPF(m) \vdash_{\mathbf{K}}^1 \beta.$$

And since $\beta \in B$ was an arbitrary conjunct of α , this means that we can also derive $\alpha = \bigwedge_{\beta \in B} \beta$ from the true propositional facts of m , as required.

Example 7.50 Similar axiomatizations can be given to *monotone* and *graded* one-step logic.

(1) The derivation system **M**, consisting of the monotonicity rule for the modality $\Box$ of monotone modal logic:

$$\mathbf{M} := \{a \rightarrow b / \Box a \rightarrow \Box b\}$$

is one-step and complete.

(2) We write $\Diamond^k$ for the counting modality $\heartsuit_{\lambda \geq k}$ associated with the predicate lifting $\lambda^{\geq k}$ for the bag functor, cf. Example 6.7. One may show that the following provides a one-step sound and complete axiomatisation for the signature consisting of all these modalities:

- a. $\Diamond^{n+1}a \rightarrow \Diamond^n a$
- b. $\Box^1(a \rightarrow b) \rightarrow (\Diamond^n a \rightarrow \Diamond^n b)$
- c. $\neg \Diamond^1(a \wedge b) \wedge \Diamond^{k_1!}a \wedge \Diamond^{k_2!}b \rightarrow \Diamond^{(k_1+k_2)!}(a \vee b)$
- d. $\Box^1 \top$

Here we abbreviate $\Box^k \pi := \neg \Diamond^k \neg \pi$ and $\Diamond^{k!} \pi \equiv \Diamond^k \pi \wedge \neg \Diamond^{k+1} \pi$.

Finally, as a Hilbert-style analog of Theorem 7.39 we have the following result. It implies that there always exists *some* one-step complete Hilbert-style derivation system; and thus, a one-step complete Hilbert-style axiomatisation, for every set functor T and for every modal signature for T .

Theorem 7.51 *Let Λ be a modal signature for a set functor T . Then the set of all one-step sound one-step Hilbert-style derivation rules is in fact a one-step complete Hilbert-style derivation system for Λ and T .*

8 Soundness and completeness

In this chapter we will see how to find sound and complete derivation systems for the formulas in ML_Λ that are *valid* in the class of all coalgebras (of the appropriate type). In a slogan, we will show that

completeness of coalgebraic logic is determined at the one-step level,

in the sense that we may transform any one-step sound and complete derivation system $\mathbf{D}$, just by adding some appropriate propositional rules, into a system $\mathbf{D}^+$ that is sound and complete for the validities in the full language ML_Λ . We will prove this result both in the setting of sequent calculi and that of Hilbert-style derivation systems.

Throughout this chapter we fix a set functor T , a modal signature Λ for T , and a countably infinite set A of variables. For convenience we will assume that T is standard.

Remark 8.1 We also assume that T is not the trivial functor $K_\emptyset$ which maps every set to the empty set, and every function to the empty map (which is the identity arrow on the empty set). It can be shown that $K_\emptyset$ admits only the empty coalgebra, and for this reason any coalgebraic modal logic for this functor will be trivial.

8.1 Sequent calculi

In this section we will see, in the framework of sequent calculi, how we can transfer soundness and completeness of a derivation system from the one-step level to the level of the full coalgebraic modal language.

8.1.1 Sequent systems

To recall the setting, given a one-step derivation system $\mathbf{G}$, we obtain a system $\mathbf{G}^+$ by adding some (fixed) propositional rules for classical logic. For simplicity we restrict the propositional connectives to the (expressively complete) set $\{\top, \neg, \wedge\}$.

Definition 8.2 Consider the following propositional (axioms and) rules:

$$\frac{}{a, \neg a} \text{Ax0} \quad \frac{}{\top} \text{Ax1} \quad \frac{a}{\neg \neg a} \text{R}_{\neg \neg} \quad \frac{a \quad b}{a \wedge b} \text{R}_{\wedge} \quad \frac{\neg a, \neg b}{\neg(a \wedge b)} \text{R}_{\neg \wedge}$$

Given a one-step derivation system $\mathbf{G}$, we define $\mathbf{G}^+$ as its extension with the above propositional rules. $\triangleleft$

We define ML_Λ -*instances* of a derivation rule in the obvious way. In particular, instances of the above propositional rules look as follows:

$$\frac{}{\Gamma, \varphi, \neg \varphi} \text{Ax0} \quad \frac{}{\Gamma, \top} \text{Ax1} \quad \frac{\Gamma, \varphi}{\Gamma, \neg \neg \varphi} \text{R}_{\neg \neg} \quad \frac{\Gamma, \varphi \quad \Gamma, \psi}{\Gamma, \varphi \wedge \psi} \text{R}_{\wedge} \quad \frac{\Gamma, \neg \varphi, \neg \psi}{\Gamma, \neg(\varphi \wedge \psi)} \text{R}_{\neg \wedge}$$

with Γ an arbitrary set of ML_Λ -formulas. In the case of a one-step rule $\mathbf{d} = \Delta_1, \dots, \Delta_{n-1} / \Delta_0$, any instance of $\mathbf{d}$ can be represented as a rule

$$\frac{\Delta_1[\tau] \quad \dots \quad \Delta_{n-1}[\tau]}{\Gamma, \Delta_0[\tau]} \mathbf{d}$$

where $\tau : A \rightarrow \text{ML}_\Lambda$ is a substitution, and $\Gamma \subseteq \text{ML}_\Lambda$ is a set of context formulas. As in the previous chapter, the difference between instances of propositional and one-step rules is that in the case of a one-step rule the context only appears in the conclusion, not in the premises.

Formally we will define rule instances as follows.

Definition 8.3 An *instance* of a derivation rule $\mathbf{d} \in \mathbf{G}^+$ is a triple $(\Theta, \mathbf{d}, \tau)$ such that $\text{conc}(\mathbf{d})[\tau] \subseteq \Theta$. The *context* of such a rule instance is the set $\Theta \setminus \text{conc}(\mathbf{d})[\tau]$. $\triangleleft$

Example 8.4 We look at standard modal logic (that is, $T = P$), taking the diamond operator $\Diamond$ as the single modality. Let $\mathbf{M}_\Diamond$ be the derivation system consisting of all modal derivation rules of the form

$$\frac{\neg a_0, a_1, \dots, a_n}{\neg \Diamond a_0, \Diamond a_1, \dots, \Diamond a_n} \mathbf{m}_n$$

In Example 7.38 we saw that this system is one-step complete. It will follow from the results in this section that the system $\mathbf{M}_\Diamond^+$ is a sound and complete derivation system for the set of all valid modal formulas in this language.

The notion of a $\mathbf{G}^+$ -*derivation* and its *assumptions* are defined in the obvious way. A $\mathbf{G}^+$ -*proof* is a finite $\mathbf{G}^+$ -derivation without assumptions; that is, all leaves are labelled with axioms. For completeness we provide the details below.

Definition 8.5 Let $\mathbf{G}$ be a one-step derivation system. A *derivation* in $\mathbf{G}^+$ is a structure (W, C, L, r) , where (W, C, r) is a well-founded tree with root r , and L is a map *labelling* every node w of W with an ML_Λ -sequent $L(w)$, in such a way that, for every inner node $t \in W$, the pair $(\{L(u) \mid u \in C(t)\}, L(t))$ is an instance of a derivation rule in $\mathbf{G}^+$.

Given an $\mathbf{G}^+$ -derivation $\mathcal{D} = (W, C, L, r)$, we call a sequent Σ an *assumption* of $\mathcal{D}$ if $\Sigma = L(t)$ for some *leaf* t , but Σ is not an instance of an axiom; we denote the set of assumptions of $\mathcal{D}$ by $\text{Ass}(\mathcal{D})$. We refer to the sequent $L(r)$ labelling the root of the tree as the *result* of $\mathcal{D}$.

A sequent $\Sigma \in 1\text{ML}_\Lambda(A)$ is *derivable* from a collection $\mathcal{A}$ of sequents in a derivation system $\mathbf{G}^+$, notation: $\mathcal{A} \vdash_{\mathbf{G}^+} \Sigma$ if Σ is the result of an $\mathbf{G}^+$ -derivation $\mathcal{D}$ with $\text{Ass}(\mathcal{D}) \subseteq \mathcal{A}$. If $\emptyset \vdash_{\mathbf{G}^+} \Sigma$ we simply write $\vdash_{\mathbf{G}^+} \Sigma$ and we say that Σ is *derivable*. $\triangleleft$

Example 8.6 Here is an example of an $\mathbf{M}_\Diamond^+$ -derivation:

$$\begin{array}{c}
\frac{}{p, \neg p, \neg \neg q} \text{Ax0} \quad \frac{}{\neg q, \neg p, \neg \neg q} \text{Ax0} \quad \frac{}{\top} \text{Ax1} \\
\frac{}{p \wedge \neg q, \neg p, \neg \neg q} R_\wedge \quad \frac{}{\neg \neg \top} R_{\neg \neg} \quad \frac{}{\neg \Diamond \neg \top} m_0 \\
\frac{}{\Diamond(p \wedge \neg q), \Diamond \neg p, \neg \Diamond \neg q, r} m_2 \quad \frac{}{\Diamond(p \wedge \neg q), \Diamond \neg p, \neg \Diamond \Diamond \neg \top, r} m_0 \\
\frac{}{\Diamond(p \wedge \neg q), \Diamond \neg p, \neg \Diamond \neg q \wedge \neg \Diamond \Diamond \neg \top, r} R_{\neg \neg} \\
\frac{}{\neg \neg \Diamond(p \wedge \neg q), \Diamond \neg p, \neg \Diamond \neg q \wedge \neg \Diamond \Diamond \neg \top, r} R_{\neg \neg} \\
\frac{}{\neg \Diamond \neg \Diamond(p \wedge \neg q), \Diamond \Diamond \neg p, \Diamond(\neg \Diamond \neg q \wedge \neg \Diamond \Diamond \neg \top), r, \Diamond r} m_3 \\
\frac{}{\neg \Diamond \neg \Diamond(p \wedge \neg q), \neg \Diamond \Diamond \neg p, \Diamond(\neg \Diamond \neg q \wedge \neg \Diamond \Diamond \neg \top), r, \Diamond r} R_{\neg \neg} \\
\frac{}{\neg(\Diamond \neg \Diamond(p \wedge \neg q) \wedge \neg \Diamond \Diamond \neg p), \Diamond(\neg \Diamond \neg q \wedge \neg \Diamond \Diamond \neg \top), r, \Diamond r} R_{\neg \wedge}
\end{array}$$

The main result of this section is the following. Recall that we read sequents *disjunctively*; hence, a sequent Σ is *valid* if the formula $\bigvee \Sigma$ is valid (that is, true at every state in every model).

Theorem 8.7 (Soundness and Completeness) *Let $\mathbf{G}$ be a set of modal derivation rules which is one-step sound and complete. Then the system $\mathbf{G}^+$ is sound and complete for the set of valid ML_Λ -sequents. That is, for any sequent $\Sigma \in \text{Seq}(\text{ML}_\Lambda)$ the following equivalence holds:*

$$\Sigma \text{ is valid iff it is derivable in } \mathbf{G}^+. \quad (68)$$

8.1.2 The derivability game

We will prove Theorem 8.7 by *game-theoretic* means. That is, with $\mathbf{G}$ we will associate a *derivability game* $\mathcal{D}_{\mathbf{G}}@ \Sigma$, with two players: *Prover* (P , female) and *Builder* (B , male). The aim of Prover in this game is to find a derivation for Σ , while Builder tries to refute Σ by constructing a *countermodel*, that is, a pointed model where all formulas in Σ are false. The derivability game is played on a board that consists of two kinds of positions:

- Basic positions are given as ML_Λ -sequents, and they are *owned* by Prover. At such a position she has to pick an instance of a rule in $\mathbf{G}^+$ that is applicable to the sequent.
- Secondary positions are given as ML_Λ -instances of derivation rules in $\mathbf{G}^+$ and these positions are *owned* by Builder. At a position of this kind, he has to choose a premise of the rule instance, and in that way he picks the next basic position of the match.

A match of the game $\mathcal{D}_{\mathbf{G}}@ \Sigma$ thus proceeds by Prover and Builder successively moving a token from one position to the next, starting from the position Σ . Such a match continues until one of the players gets stuck.

Formally, the game consists of a collection of positions, that provide the vertices of a directed graph (Pos, E) . The idea of the game is that, starting from the initial position Σ , the players move a token around the graph, from one position to another, following the edge relation E . Each position is *owned* by one of the players. This ownership relation indicates

Position $p \in \text{Pos}$	Owner $O(p)$	Admissible moves $E(p) \subseteq \text{Pos}$
Θ	P	$\{(\Theta, d, \tau) \mid d \in \mathbf{G}^+, \tau \in \text{ML}_\Lambda^A, \text{conc}(d)[\tau] \subseteq \Theta\}$
(Θ, d, τ) , d propositional	B	$\{\Delta[\tau] \cup (\Theta \setminus \text{conc}(d)[\tau]) \mid \Delta \in \text{Prem}(d)\}$
(Θ, d, τ) , d one-step	B	$\{\Delta[\tau] \mid \Delta \in \text{Prem}(d)\}$

Table 2: Derivability game for a derivation system $\mathbf{G}^+$

that, when the token is situated at a position, it is the task of its owner to move the token to a neighbouring position. Here are the formal details of the nature of positions and the edge relation linking them.

- Positions of the form Θ , for some ML_Λ -sequent Θ , are called basic and belong to Prover. She needs to pick a rule $d \in \mathbf{G}^+$ and a substitution $\tau : A \rightarrow \text{ML}_\Lambda$ such that $\text{conc}(d)[\tau] \subseteq \Theta$. Doing so she moves the token to the position (Θ, d, τ) .
- Positions of the form (Θ, d, τ) , with $\text{conc}(d)[\tau] \subseteq \Theta$, belong to Builder. We think of Θ as the conclusion of the rule instance induced by d and τ and we partition the sequent as $\Theta = \text{conc}(d)[\tau] \uplus \Gamma$, where $\text{conc}(d)[\tau]$ and Γ consist of, respectively, the *principal* and *context* formulas in Θ . It is now Builder's task to choose a premise Δ of d ; in doing so he picks the next position Θ' in the match as follows:

$$\Theta' := \begin{cases} \Delta[\tau] \cup \Gamma & \text{if } d \text{ is a propositional rule,} \\ \Delta[\tau] & \text{if } d \text{ is a modal rule.} \end{cases}$$

If the token arrives at a position without neighbours (that is, a node $p \in \text{Pos}$ such that $E(p) = \emptyset$), we say that the position's owner *gets stuck*, and when this happens the player immediately loses the match of the game. Both players might find themselves in this situation:

- Prover gets stuck at a sequent position if there is no rule applicable to the sequent. Examples of such sequents include $\{p, \neg q\}$, $\{\neg \top\}$, and (in the case of the derivation system $\mathbf{M}_\Diamond^+$ of Example 8.4) $\{p, \Diamond p, \Diamond \neg p, \Diamond \neg q\}$ or $\{\Diamond \top\}$.
- Builder gets stuck at a position of the form (Θ, d, τ) if d is a rule without premises, that is, if Θ is an instance of an axiom.

Definition 8.8 The *derivability game* $\mathcal{D}_\mathbf{G}$ is defined as the two-player game (Pos, O, E) of which the game graph (Pos, E) and owner function $O : \text{Pos} \rightarrow \{P, B\}$ are given in Table 2. For each player $Q \in \{P, B\}$, we denote by $\text{Pos}_Q := O^{-1}(Q)$ the set of positions they own. Given an ML_Λ -sequent Σ , we let $\mathcal{D}_\mathbf{G}@\Sigma$ denote the version of the game $\mathcal{D}_\mathbf{G}$ that is initialised at Σ . $\triangleleft$

Before we show how this derivability game can be used to prove soundness and completeness of the derivation system $\mathbf{G}^+$, we prove two fundamental results about the game itself.

First we show that $\mathcal{D}_{\mathbf{G}}$ -matches always end after finitely many moves; that is, all matches are finite. A *match* of this game is a sequence of positions that arise from Prover and Builder successively moving a token from one position to another, following the rules of the game. Formally then, a match is nothing but a path through the game graph $(\mathbf{Pos}, E)$. Such a match π is *full* if it is either infinite or it is finite and its final position $\text{last}(\pi)$ has no successors in the game graph. In the latter case, the owner of this final position got stuck and so their opponent is declared to be the winner of π . Matches that are not full are called *partial*. Matches of the initialised game $\mathcal{D}_{\mathbf{G}}@ \Sigma$ must start at Σ .

To establish that all $\mathcal{D}_{\mathbf{G}}$ -matches are in fact finite, we need some complexity measures for formulas and sequents.

Definition 8.9 The *modal depth* of ML_{Λ} -formulas is given by the following formula induction:

$$\begin{aligned} \text{md}(\varphi) &:= 0 && \text{if } \varphi \text{ is atomic} \\ \text{md}(\neg\varphi) &:= \text{md}(\varphi) \\ \text{md}(\varphi_0 \wedge \varphi_1) &:= \max(\text{md}(\varphi_0), \text{md}(\varphi_1)) \\ \text{md}(\heartsuit_{\lambda}(\varphi_1, \dots, \varphi_n)) &:= 1 + \max\{\text{md}(\varphi_1), \dots, \text{md}(\varphi_n)\} \end{aligned}$$

The modal depth of a sequent Φ is defined as $\text{md}(\Phi) := \max\{\text{md}(\varphi) \mid \varphi \in \Phi\}$.

We define the *propositional complexity* of ML_{Λ} -formulas by the following induction:

$$\begin{aligned} \text{pc}(\varphi) &:= 0 && \text{if } \varphi \text{ is atomic} \\ \text{pc}(\neg\varphi) &:= 1 + \text{pc}(\varphi) \\ \text{pc}(\varphi_0 \wedge \varphi_1) &:= 1 + \text{pc}(\varphi_0) + \text{pc}(\varphi_1) \\ \text{pc}(\heartsuit_{\lambda}(\varphi_1, \dots, \varphi_n)) &:= 0 \end{aligned}$$

The propositional complexity of a sequent Φ is given as $\text{pc}(\Phi) := \sum_{\varphi \in \Phi} \text{pc}(\varphi)$. $\triangleleft$

Proposition 8.10 $\mathcal{D}_{\mathbf{G}}$ is a finite game; that is, all matches of $\mathcal{D}_{\mathbf{G}}$ are finite.

Proof. We leave it as an exercise for the reader to prove this result, using the complexity measures of Definition 8.9. QED

As our second result on the derivability game we will show that it enjoys the property of *positional determinacy*. To explain this notion we need to discuss one of the key concepts in game theory, viz., that of a player having a *winning strategy* in an (initialised) game. Intuitively, a winning strategy for player $Q \in \{P, B\}$ in $\mathcal{D}_{\mathbf{G}}@ \Sigma$ is a way of playing for Q which guarantees, no matter the response of Q 's opponent $\overline{Q}$, that Q will win any match starting from Σ , as long as they stick to the given strategy. The winning region for Q in a game consists of the set of those positions for which Q has a winning strategy.

► say something about formalising these concepts in general

In the case of finite games, however, we may always assume that winning strategies are *positional*.

Definition 8.11 A *positional strategy* for player $Q \in \{P, B\}$ in $\mathcal{D}_{\mathbf{G}}$ is a partial map $f : \text{Pos}_Q \xrightarrow{\circ} \text{Pos}$ such that $f(p) \in E(p)$, for every position p in the domain of f . A match $\pi = p_0 \cdots p_n$ is *guided* by a positional strategy f for Q if $p_{i+1} = f(p_i)$ for every $i < n$ such that $p_i \in \text{Pos}_Q$.

A positional strategy is *winning for Q* at position $p \in \text{Pos}$ if for every f -guided match π such that $\text{first}(\pi) = p$ and $\text{last}(\pi) \in \text{Pos}_Q$ we have $\text{last}(\pi) \in \text{Dom}(f)$. A positional strategy is *winning* for a region $P \subseteq \text{Pos}$ if $P \subseteq \text{Dom}(f)$ and f is winning for every position $p \in P$. $\triangleleft$

Observe that a positional strategy f being winning for a player at first glance only ensures that the player will never get stuck when playing f ; but since in this game loosing means getting stuck, simply surviving actually means being guaranteed to win.

The main property of the derivability game that we shall need below is the following.

Theorem 8.12 (Positional Determinacy) *The set Pos of positions in the game $\mathcal{D}_{\mathbf{G}}$ can be partitioned as*

$$\text{Pos} = W_P \uplus W_B,$$

in such a way that each player Q has a positional strategy which is winning for the region W_Q .

► Proof (sketch) to be supplied

It is not hard to see that the partition mentioned in Theorem 8.12 is in fact unique. This inspires and justifies the following definition.

Definition 8.13 We let Win_P and Win_B denote the sets given in Theorem 8.12. For a given player $Q \in \{P, B\}$, we refer to Win_Q as the *winning region of Q* . $\triangleleft$

Note that the winning region for player Q may contain both positions that are owned by Q and positions that are owned by Q 's opponent.

8.1.3 Soundness and completeness

With the definition and basic properties of the derivability game in place, we will now show how $\mathcal{D}_{\mathbf{G}}$ can be used to prove the soundness and completeness of the derivation system $\mathbf{G}^+$. The intuitions underlying the proof are quite straightforward: a winning strategy for Prover in the initialised game $\mathcal{D}_{\mathbf{G}}@ \Sigma$ corresponds to a *proof* of Σ in $\mathbf{G}^+$, while a winning strategy for Builder can be seen as a *countermodel* for Σ , or at least as a blueprint for constructing such a model. Soundness and completeness of $\mathbf{G}^+$ then follow from the determinacy of the derivability game: since either Prover or Builder must have a winning strategy in $\mathcal{D}_{\mathbf{G}}@ \Sigma$, we are guaranteed to find either a proof or a countermodel for Σ . Here is a more formal proof.

Proof of Theorem 8.7. Let $\mathbf{G}$ be some one-step sequent system which is one-step sound and complete, and let $\Sigma \in \text{Seq}(\text{ML}_{\Lambda})$ be some sequent. Consider the following chain of

equivalent statements:

Σ is derivable in $\mathbf{G}^+$	iff $\Sigma \in \text{Win}_P(\mathcal{D}_{\mathbf{G}})$	(Proposition 8.14)
	iff $\Sigma \notin \text{Win}_B(\mathcal{D}_{\mathbf{G}})$	(Theorem 8.12)
	iff Σ does not have a countermodel	(Proposition 8.15)
	iff Σ is valid	(obvious)

This obviously proves the Theorem.

QED

Thus, all that is left is to prove the two Propositions below. The proof of the first Proposition is left to the reader.

Proposition 8.14 *Let Σ be some ML_Λ -sequent. Then Prover has a winning strategy in $\mathcal{D}_{\mathbf{G}}@ \Sigma$ iff Σ has a $\mathbf{G}^+$ -proof.*

The heart of the completeness proof is the Proposition below.

Proposition 8.15 *Let Σ be some ML_Λ -sequent. Then Builder has a winning strategy in $\mathcal{D}_{\mathbf{G}}@ \Sigma$ iff Σ is refutable.*

Proof. For the direction from left to right, we will in fact show something stronger than the statement in the Proposition: we will prove the existence of a *single* model that refutes *every* sequent in Builder's winning region in $\mathcal{D}_{\mathbf{G}}@ \Sigma$. The domain of this canonical countermodel consists precisely of these sequents themselves:

$$S := \text{Seq}(\text{ML}_\Lambda) \cap \text{Win}_B(\mathcal{D}_{\mathbf{G}}).$$

For the definition of the coalgebra map and the valuation of this model we fix some *positional winning strategy* f for Builder in $\mathcal{D}_{\mathbf{G}}$. Recall that such a strategy exists by Theorem 8.12, and that it is a function that maps any position of the form $(\Theta, \mathbf{d}, \tau)$ to a premise of the rule instance that the position represents. Furthermore, we will confine attention to f -guided matches where Prover plays 'supportively'. To explain this we need some preparations.

First, we call a sequent Θ *propositionally simple* if each of its members has propositional complexity zero or one. In other words, a sequent is propositionally simple if it contains no formulas of the form $\neg\neg\varphi$, $\varphi \wedge \psi$ or $\neg(\varphi \wedge \psi)$; that is, all of its elements are atomic or modal formulas, or negations of such formulas. Note that only modal rules may be applied to a propositionally simple sequent (unless it is axiomatic).

For the definition of the coalgebra map $\sigma_f : S \rightarrow TS$ and the valuation V_f we confine our attention to f -guided matches where Prover plays 'supportively'. That is, we fix some strict linear order $\sqsubset$ on formulas that is compatible with the subformula order ($\varphi \triangleleft \psi$ implies $\varphi \sqsubset \psi$), and require that

- Prover postpones playing modal rules until no propositional rules are available; and
- when playing a propositional rule, Prover picks the principal formula that is *maximal* with respect to $\sqsubset$.

Claim 1 Let Θ be a sequent in S . If Prover plays supportively, then there is a *unique* f -guided partial match π such that $\text{first}(\pi) = \Theta$, $\text{last}(\pi)$ is propositionally simple, and P only plays propositional rules during π .

Given a sequent state $\Theta \in S$, we let π_Θ denote the match given by the above Claim, and we define

$$\begin{aligned}\Theta^\ell &:= \text{last}(\pi_\Theta) \\ \Theta^+ &:= \bigcup \{ \Phi \mid \Phi \text{ occurs on } \pi_\Theta \}.\end{aligned}$$

In words, Θ^+ consists of all formulas that belong to some sequent on the path π_Θ .

Claim 2 For any $\Theta \in S$, the set Θ^+ is saturated; that is, for all formulas φ, ψ we have:

- (1) it is not the case that $\varphi, \neg\varphi \in \Theta^+$ or that $\top \in \Theta$;
- (2) if $\neg\neg\varphi \in \Theta^+$, then also $\varphi \in \Theta^+$;
- (3) if $\varphi \wedge \psi \in \Theta^+$, then $\varphi \in \Theta^+$ or $\psi \in \Theta^+$;
- (4) if $\neg(\varphi \wedge \psi) \in \Theta^+$, then also $\neg\varphi, \neg\psi \in \Theta^+$.

In addition we have that

- (5) $\Theta^\ell \subseteq \Theta^+$;
- (6) $\varphi \in \Theta^\ell \iff \varphi \in \Theta^+$, for all φ of the form $\pm p$ or $\pm\heartsuit_\lambda(\psi_0, \dots, \psi_{n-1})$.

The definition of the valuation V_f is straightforward:

$$V_f(p) := \{ \Theta \in S \mid \neg p \in \Theta^+ \}.$$

For the coalgebra map σ_f we have to do more work; we will define σ_f layer by layer, where the layers of S are given by the modal depth of sequents. Recall the notion of the *modal depth* $\text{md}(\Theta)$ of a sequent Θ from Definition 8.9.

Note that the sequents that we subsequently encounter in a match of the derivability game have *decreasing* modal depth. In particular, if Prover plays a modal rule at a sequent position Θ , then the next sequent of the match will have a *strictly smaller* modal depth than Θ : the modal rule peels off one layer of modalities. If she plays a propositional rule, the modal depth may decrease or stay the same, but it will not increase.

The layers of S are then given by the following definition:

$$S_n := \{ \Theta \in S \mid \text{md}(\Theta) \leq n \}.$$

By standardness of T have $TS_m \subseteq TS_n$ if $m < n$. The following observation is the pivotal claim in the proof.

Claim 3 Assume that, for some $n \in \omega$, there is a map $\sigma_n : S_n \rightarrow TS_n$ such that in the resulting coalgebra model $\mathbb{S}_n = (S_n, \sigma_n, V_f \upharpoonright_{S_n})$ every state $\Phi \in S_n$ refutes the sequent Φ^+ . Then we can extend σ_n to a map $\sigma_{n+1} : S_{n+1} \rightarrow TS_n$ satisfying the following two conditions:

- (1) For every $\Theta \in S_{n+1} \setminus S_n$ and every formula $\heartsuit_\lambda(\psi_0, \dots, \psi_{n-1})$ we have, writing $\tau := \sigma_{n+1}(\Theta)$:

$$\begin{aligned}\heartsuit_\lambda(\psi_0, \dots, \psi_{n-1}) \in \Theta^+ &\quad \text{implies} \quad \tau \notin \lambda_{S_n}(\llbracket \psi_0 \rrbracket^{\mathbb{S}_n}, \dots, \llbracket \psi_{n-1} \rrbracket^{\mathbb{S}_n}) \\ \neg\heartsuit_\lambda(\psi_0, \dots, \psi_{n-1}) \in \Theta^+ &\quad \text{implies} \quad \tau \in \lambda_{S_n}(\llbracket \psi_0 \rrbracket^{\mathbb{S}_n}, \dots, \llbracket \psi_{n-1} \rrbracket^{\mathbb{S}_n}).\end{aligned}\tag{69}$$

- (2) In the resulting coalgebra model $\mathbb{S}_{n+1} = (S_{n+1}, \sigma_{n+1}, V_f \upharpoonright_{S_{n+1}})$ every state $\Phi \in S_{n+1}$ refutes the sequent Φ^+ .

PROOF OF CLAIM For part 1, fix some $\Theta \in S_{n+1} \setminus S_n$, and assume for contradiction that there is *no* object $\tau \in TS_n$ meeting the conditions in (69). To avoid notational clutter we assume that all predicate liftings in Λ are unary – this is not without loss of generality, but the general case is not essentially harder. Introduce a variable a_ψ for every formula ψ such that either $\heartsuit_\lambda \psi \in \Theta^+$ or $\neg \heartsuit_\lambda \psi \in \Theta^+$, for some $\lambda \in \Lambda$; let A be the set of all these variables; and let $\nu : A \rightarrow \text{ML}_\Lambda$ be the natural substitution given by $\nu(a_\psi) := \psi$. Furthermore, define

$$\Gamma := \{\heartsuit_\lambda a_\psi \mid \heartsuit_\lambda \psi \in \Theta^+\} \cup \{\neg \heartsuit_\lambda a_\psi \mid \neg \heartsuit_\lambda \psi \in \Theta^+\},$$

then we obviously have $\Gamma[\nu] \subseteq \Theta^+$.

By our assumption on the non-existence of an element $\tau \in TS_n$ satisfying (69) we can find, for every $\tau \in TS_n$, a formula $\varphi_\tau \in \Theta^+$ such that either $\varphi_\tau = \heartsuit_\lambda \psi$ and $\tau \in \lambda_{S_n}(\llbracket \psi \rrbracket^{\mathbb{S}_n})$ or $\varphi_\tau = \neg \heartsuit_\lambda \psi$ and $\tau \notin \lambda_{S_n}(\llbracket \psi \rrbracket^{\mathbb{S}_n})$. Hence, if we define $m : S_n \rightarrow PA$ by

$$m : \Phi \mapsto \{a_\psi \in A \mid \mathbb{S}_n, \Phi \Vdash \psi\},$$

we obtain, for each $\tau \in TS_n$, a one-step formula α in Γ of the form $\alpha = \heartsuit_\lambda \psi$ (if $\heartsuit_\lambda \psi \in \Theta^+$) or $\alpha = \neg \heartsuit_\lambda a_\psi$ (if $\neg \heartsuit_\lambda \psi \in \Theta^+$) such that $\tau \in \llbracket \alpha \rrbracket_m^1$. In other words we find that $\llbracket \Gamma \rrbracket_m^1 = TS_n$.

But then by one-step completeness there must be some modal rule $\mathbf{d} = \Delta_1 \cdots \Delta_{n-1} / \Gamma_0$ in $\mathbf{G}$ and a renaming $\tau : A \rightarrow A$ such that $\Gamma_0[\tau] \subseteq \Gamma$ and $\llbracket \Delta_i[\tau] \rrbracket_m^0 = S$, for all i .

On the other hand, by construction of Γ we have $\Gamma_0[\nu \circ \tau] = \Gamma_0[\tau][\nu] \subseteq \Gamma[\nu] \subseteq \Theta$. But then by the observation in Claim 2(6) we have $\Gamma_0[\nu \circ \tau] \subseteq \Theta^\ell$, so that the triple $(\Gamma_0, \mathbf{d}, \nu \circ \tau)$ is a legitimate move for Prover at position Θ^ℓ . Then for each i , the sequent $\Delta_i[\nu \circ \tau]$ is a premise of this rule instance, and so one of these, say, $\Phi := \Delta_i[\nu \circ \tau]$ is the response picked by Builder's winning strategy f if Prover makes this move $(\Gamma_0, \mathbf{d}, \nu \circ \tau)$ indeed. Note that $\Phi = \Delta_i[\tau][\nu]$ implies

$$\Delta_i = \{a_\varphi \mid \varphi \in \Phi\} \cup \{\neg a_\varphi \mid \neg \varphi \in \Phi\},$$

and that $\text{md}(\Phi) = \text{md}(\Theta) - 1 = n$.

In other words, we have $\Phi \in S_n$, and so Φ , as a state in $\mathbb{S}_n$, refutes Φ^+ by assumption. That is, we find that $\mathbb{S}_n, \Phi \Vdash \varphi$ for no $\varphi \in \Phi^+$, which by definition of m means that $\Phi \notin \llbracket \Delta_i[\tau] \rrbracket_m^0$. This, however, contradicts the fact that $\llbracket \Delta_i[\tau] \rrbracket_m^0 = S$.

For part 2 of the claim take an arbitrary sequent $\Theta \in S_{n+1}$, and distinguish cases. If Θ belongs to S_n , note that since σ_{n+1} extends σ_n we have $\sigma_{n+1}(\Theta) = \sigma_n(\Theta) \in TS_n$. In other words, (S_n, σ_n) is a subcoalgebra of (S_{n+1}, σ_{n+1}) and so we have that $\mathbb{S}_{n+1}, \Theta \Vdash \varphi$ iff $\mathbb{S}_n, \Theta \Vdash \varphi$, for *any* formula φ . In particular this means that Θ , as a state in $\mathbb{S}_{n+1}$, refutes Θ^+ , by the assumption that it does so as a state in $\mathbb{S}_n$.

In case $\Theta \in S_{n+1} \setminus S_n$ it suffices to show that any formula $\varphi \in \text{PL}(\Lambda(\mathbf{Sfor}_0(\Theta^+)))$ satisfies the following conditions:

$$\begin{array}{ll} \varphi \in \Theta^+ & \text{implies } \mathbb{S}_{n+1}, \Theta \not\Vdash \varphi \\ \neg \varphi \in \Theta^+ & \text{implies } \mathbb{S}_{n+1}, \Theta \Vdash \varphi. \end{array}$$

We prove this statement by induction on the propositional complexity of φ .

In the base case of this induction φ is either atomic or of the form $\varphi = \heartsuit_\lambda(\vec{\psi})$. We leave the first case as an exercise for the reader. In the second case the above statements readily follow from part 1 of the Claim and the definition of $\mathbb{S}_{n+1}$.

In the induction step, we have two cases to consider. We leave the case where φ is of the form $\varphi = \neg\varphi'$ as an exercise for the reader. In the case where φ is of the form $\varphi = \varphi_0 \wedge \varphi_1$ we reason as follows. If $\varphi \in \Theta^+$ then by Claim 2 we find φ_i in Θ^+ for some $i \in \{0, 1\}$, so that by the induction hypothesis we obtain $\mathbb{S}_{n+1}, \Phi \not\models \varphi_i$ for this i . But then obviously we do not have $\mathbb{S}_{n+1}, \Phi \models \varphi$ either. On the other hand, if $\neg\varphi \in \Theta^+$ then by Claim 2 we find both formulas $\neg\varphi_0, \neg\varphi_1$ in Θ^+ . By the induction hypothesis this means that $\mathbb{S}, \Theta \models \varphi_i$ for both $i = 0, 1$, and clearly this implies that $\mathbb{S}, \Theta \models \varphi_0 \wedge \varphi_1$, as required. $\triangleleft$

Based on Claim 3 we can now easily prove the following.

Claim 4 For every $n \in \omega$ there is a map $\sigma_n : S_n \rightarrow TS_n$ such that, in the resulting coalgebra model $\mathbb{S}_n = (S_n, \sigma_n, V_f)$ we have that Φ refutes Φ^+ , for all $\Phi \in S_n$.

PROOF OF CLAIM We prove the claim by a straightforward induction. For the base case of this induction, where $n = 0$, observe that sequents in S_0 only contain propositional formulas. Given such a sequent Θ , it easily follows from Claim 2 that both Θ and Θ^+ are satisfiable sets of propositional formulas, and that, by the definition of the valuation V_f , Θ refutes Θ^+ . However, we still need to define a coalgebra map $\sigma_0 : S_0 \rightarrow TS_0$, but for $\Theta \in S_0$ we may define $\sigma_0(\Theta)$ to be any member¹² of the set $T\{\Theta\}$. (Note that the definition of $\sigma_0(\Theta)$ has no impact on the truth of propositional formulas at Θ .)

In the inductive case we assume the existence of a map $\sigma_n : S_n \rightarrow TS_n$ satisfying the conditions of the claim, and we simply use the previous claim to find an appropriate map $\sigma_{n+1} : S_{n+1} \rightarrow TS_{n+1}$. $\triangleleft$

It is then a straightforward consequence of Claim 4 that the initial sequent Σ is refuted in the model $\mathbb{S}$. This finishes the proof of the left-to-right direction of the equivalence stated in the Proposition.

The opposite, right-to-left direction of the proof is left as an exercise for the reader.

QED

8.2 Hilbert calculi

► Section to be added later

¹²In case this set is empty we are dealing with the constant functor $K_\emptyset$ which maps every set S to the empty set. This functor has only one coalgebra, namely, the empty one $(\emptyset, \emptyset)$. It follows that in this case *every sequent* is valid.

A Appendix: Basic mathematical definitions

In this second appendix we fix notation and terminology for some basic mathematical concepts. First we consider sets, functions and relations.

A binary relation between two sets X and Y is nothing but a set $R \subseteq X \times Y$ of pairs from X and Y , respectively.

Definition A.1 Let $f : X \rightarrow Y$ be a function. We let $\text{Gr}f := \{(x, y) \in X \times Y \mid y = fx\}$ denote the *graph* of f , and define $f[X] := \{fx \mid x \in X\}$. $\triangleleft$

Definition A.2 Given a relation $R \subseteq X \times Y$, we denote the *domain* $\text{Dom}(R) \subseteq X$ and *range* $\text{Ran}(R) \subseteq Y$ of R by the followings sets:

$$\begin{aligned}\text{Dom}(R) &:= \{x \in X \mid (x, y) \in R \text{ for some } y \in Y\} \\ \text{Ran}(R) &:= \{y \in Y \mid (x, y) \in R \text{ for some } x \in X\},\end{aligned}$$

respectively, and we denote by $\pi_0^R : R \rightarrow X$ and $\pi_1^R : R \rightarrow Y$ the projection maps associated with R . Given subsets $X' \subseteq X$, $Y' \subseteq Y$, the *restriction* of R to X' and Y' is given as

$$R \upharpoonright_{X' \times Y'} := R \cap (X' \times Y').$$

The *converse* of R is defined as the relation $R^\smile \subseteq Y \times X$ given by

$$R^\smile := \{(y, x) \in Y \times X \mid (x, y) \in R\}.$$

The *composition* of two relations $R \subseteq X \times Y$ and $R' \subseteq Y \times Z$ is denoted by $R; R'$ and defined as

$$R; R' := \{(x, z) \in X \times Z \mid (x, y) \in R \text{ and } (y, z) \in R', \text{ for some } y \in Y\}.$$

Finally, we let

$$\Delta_X := \{(x, x) \in X \times X \mid x \in X\}$$

denote the *diagonal relation* on a set X . $\triangleleft$

Proposition A.3 The transformations below constitute a bijection between the set $P(X \times Y)$ of binary relations between X and Y and the collection $(PY)^X$ of maps from X to the power set of Y :

- given $R \subseteq X \times Y$ consider $F_R : X \rightarrow PY$ given by $F_R(x) := \{y \in Y \mid (x, y) \in R\}$;
- given $F : X \rightarrow PY$, define $R_F \subseteq X \times Y$ as $R_F := \{(x, y) \in X \times Y \mid y \in F(x)\}$.

Related to this, and to the operation of taking the converse of a binary relation, is the following.

Definition A.4 Given a map $F : X \rightarrow PY$, we define the map $F^\flat : Y \rightarrow PX$ by putting,

$$F^\flat(y) := \{x \in X \mid y \in F(x)\},$$

and we call $F^\flat$ the *transpose* of F . $\triangleleft$

Note that the relations associated with, respectively, a map $F : X \rightarrow P$, and its transpose $F^b : Y \rightarrow P$ are in fact each other's converse. Observe as well that we have $F = (F^b)^b$, for every map $F : X \rightarrow PY$.

Example A.5 In case Q is a set of proposition letters or variables and S is some set of states, we tend to think of a map from Q to PS as a *valuation*, and of its transpose $F^b : S \rightarrow PQ$ as its associated *colouring* or *marking*.

Definition A.6 A *relational structure* or *graph* is a pair $\mathbb{S} = (S, R)$ such that R is a binary relation on S , that is, a set $R \subseteq S \times S$ of S -pairs. In the line of Proposition A.3 we will often think of such a relation R as a map $R : S \rightarrow PS$ given by

$$R(s) := \{t \in S \mid (s, t) \in R\}.$$

Elements of the set $R(s)$ will be called *R-successors* of s . ◁

Definition A.7 Fix a relational structure $\mathbb{S} = (S, R)$. A *path* in $\mathbb{S}$ is a sequence $\pi = (t_i)_{0 \leq i < \kappa}$, where κ is either finite or equal to ω , and we have $t_{i+1} \in R(t_i)$ for all i such that $0 \leq i < \kappa$. The *length* of $\pi = (t_i)_{0 \leq i < \kappa}$ is defined as $|\pi| := \kappa$, and we call π *proper* if $\kappa > 0$, *finite* if $0 \leq \kappa < \omega$, and *infinite* if $\kappa = \omega$. We define $\text{first}(\pi) := t_0$, and, provided π is finite, $\text{last}(\pi) := t_{\kappa-1}$. If π is finite we say it is a path *from first(π) to last(π)*. ◁

Definition A.8 A *tree* is a structure $\mathbb{T} = (T, C, r)$ such that C is a binary relation on T ; and r is the *root* of the tree, that is, r is an element of T , such that for every node $t \in T$ there is exactly one path from r to t .

Elements of trees are usually referred to as *nodes*. A *leaf* of $\mathbb{T}$ is a node $t \in T$ such that $C(t) = \emptyset$; nodes that are not leaves are called *inner* nodes. If we have $t \in C(s)$ we say that t is a *child* of s , and, vice versa, that s is the (unique) *parent* of t ; distinct children of a node are called *siblings*.

A tree is *finitely branching* if $C(t)$ is finite, for all nodes t ; and *well founded* if it has no infinite paths. ◁

It is not hard to prove that the root of a tree is completely determined by its relational structure; that is, we have $r = r'$ if both (T, C, r) and (T, C, r') are trees. For this reason we sometimes represent a tree (T, C, r) simply as the pair (T, C) .

B Appendix: The Category Set and its Functors

The theory of coalgebra is categorical in nature. In this appendix we summarize the background knowledge on category theory that is required for understanding the notes; we place a special emphasis on the category **Set** of sets and functions, since this is the base category of most of the coalgebras that we consider.

For a proper introduction to category theory, the reader is referred to standard textbooks such as Mac Lane's *Categories for the Working Mathematician*, or Awodey's *Category Theory*¹⁴ on which we based parts of this appendix.

B.1 Categories, functors and natural transformations

Definition B.1 A *category* $\mathbf{C}$ consists of a class $\text{Ob}(\mathbf{C})$ of *objects*, and for each pair of objects A, B , a family $\mathbf{C}(A, B)$ of *arrows*. If f belongs to the latter set, we write $f : A \rightarrow B$, and call A the *domain* and B the *codomain* of the arrow. The collection of arrows is endowed with some algebraic structure: for every object A of $\mathbf{C}$ there is an *identity* arrow $\text{id}_A : A \rightarrow A$, and every pair $f : A \rightarrow B$, $g : B \rightarrow C$ can be uniquely *composed* to an arrow $g \circ f : A \rightarrow C$. These operations are supposed to satisfy the associative law for composition, while the appropriate identity arrows are left- and right neutral elements. $\triangleleft$

An arrow $f : A \rightarrow B$ is an *iso* if it has an *inverse*, that is, an arrow $g : B \rightarrow A$ such that $f \circ g = \text{id}_B$ and $g \circ f = \text{id}_A$.

Example B.2 (a) We let **Set** denote the category with sets as objects and functions as arrows, with identity arrows and the composition of two arrows defined in the familiar way.

(b) The category **Rel** has the same objects as **Set**, but for the set of arrows $\text{Rel}(S', S)$ we take the collections of all binary relations between S' and S , with the identity arrows and the composition of two arrows defined in the obvious way.

(c) The *opposite* category $\mathbf{C}^{op}$ of a given category $\mathbf{C}$ has the same objects as $\mathbf{C}$, while $\mathbf{C}^{op}(A, B) = \mathbf{C}(B, A)$ for all objects A, B from $\mathbf{C}$, and the operations on arrows are defined in the obvious way.

Definition B.3 A *functor* $F : \mathbf{C} \rightarrow \mathbf{D}$ from a category $\mathbf{C}$ to a category $\mathbf{D}$ consists of an operation mapping objects and arrows of $\mathbf{C}$ to objects and arrows of $\mathbf{D}$, respectively, in such a way that $Ff : FA \rightarrow FB$ if $f : A \rightarrow B$, $F(\text{id}_A) = \text{id}_{FA}$ and $F(g \circ f) = (Fg) \circ (Ff)$ for all objects and arrows involved. A functor $F : \mathbf{C} \rightarrow \mathbf{D}^{op}$ is sometimes called a *contravariant* functor from $\mathbf{C}$ to $\mathbf{D}$. An *endofunctor* on $\mathbf{C}$ is a functor $F : \mathbf{C} \rightarrow \mathbf{C}$. $\triangleleft$

Definition B.4 Let $F, G : \mathbf{C} \rightarrow \mathbf{D}$ be two functors. A *natural transformation* $\alpha : F \rightarrow G$ consists of a family of maps $\alpha_A : FA \rightarrow GA$, indexed by the collection of objects of $\mathbf{C}$, such that $Gf \circ \alpha_A = \alpha_B \circ Ff$, for every arrow $f : A \rightarrow B$ in $\mathbf{C}$. In a diagram:

$$\begin{array}{ccccc} A & & FA & \xrightarrow{\alpha_A} & GA \\ f \downarrow & & Ff \downarrow & & \downarrow Gf \\ B & & FB & \xrightarrow{\alpha_B} & GB \end{array} \quad (73)$$

¹⁴S. Awodey, *Category Theory* (2nd edition), Oxford University Press, 2010.

◁

B.2 Set functors

Definition B.5 A *set functor* is a (covariant) endofunctor T on the category **Set**. ◁

Below we give some examples of set functors and of operations on set functors.

Example B.6 (a) Given a set C , we let K_C denote the *constant functor* which maps every set S to the set C , and every map $f : S \rightarrow S'$ to the identity map on C . The functor K_C is often simply denoted as C .

(b) The *identity functor* Id is the set functor that maps every object to itself, and similarly maps every arrow to itself.

(c) The *powerset functor* P maps any set S to its power set PS , and any function $f : S \rightarrow S'$ to the *direct image map* $Pf : PS \rightarrow PS'$ given by $Pf : A \mapsto \{fx \mid x \in A\}$. The *finitary power set functor* P_ω is defined similarly, with the difference that $P_\omega S$ only takes the *finite* subsets of S .

(d) The *contravariant powerset functor* $\check{P}$ also maps a set S to its power set $\check{P}S = PS$, but it maps a function $f : S \rightarrow S'$ to the *inverse image map* $\check{P}f : PS' \rightarrow PS$ given by $\check{P}f : X' \mapsto \{x \in S \mid fx \in X'\}$.

(e) Define the covariant set functor $N : \mathbf{Set} \rightarrow \mathbf{Set}$ as the composition of the contravariant power set with itself, $N := \check{P} \circ \check{P}$. Restricting this example somewhat, we may obtain various interesting functors. For instance, take the functor M given by $MS := \{\mathcal{U} \in NS \mid \mathcal{U} \text{ is upward closed with respect to } \subseteq\}$ and, for $f : S \rightarrow S'$, $Mf = (Nf) \upharpoonright_{MS}$ (it requires a short argument to prove that this defines a functor indeed). N and M are called the *neighbourhood* and the *monotone neighbourhood functor*, respectively.

(f) The *distribution functor* D assigns to a set S the collection $D(S)$ of (discrete) probability distributions over S , i.e., the set of all maps $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$. On arrows, D acts as follows: given a map $f : S \rightarrow S'$ and a probability distribution $\mu \in D(S)$, we define the map $(Df)\mu$ on S' by putting

$$(Df)(\mu)(s') := \sum_{s \in f^{-1}(s')} \mu(s).$$

We leave it for the reader to verify that D is indeed a set functor. The main points to check are (i) that $(Df)(\mu)$ is indeed a probability distribution on S' , for any $\mu \in D(S)$, and that (ii) $D(g \circ f) = (Dg) \circ (Df)$.

The *finitary* distribution functor D_ω is defined as the restriction of D to probability distributions that have finite *support*, that is, $D_\omega(S) := \{\mu \in DS \mid |S \setminus \mu^{-1}(0)| < \omega\}$. On functions, D_ω is defined as D .

(g) The *bag functor* B is defined analogously. Let $\mathbb{N}^\infty$ be the set $\mathbb{N} \cup \{\infty\}$ of natural numbers extended with the ‘number’ ∞ . We extend the standard addition operation on $\mathbb{N}$ to $\mathbb{N}^\infty$ by putting $n + \infty = \infty + n = \infty + \infty = \infty$ and defining the sum of infinitely many non-zero numbers to be ∞ as well.

Then we define $BS := (\mathbb{N}^\infty)^S$ as the set of *weight functions* $\mu : S \rightarrow \mathbb{N} \cup \{\infty\}$. On arrows, B acts similarly as D : given a map $f : S \rightarrow S'$ and a weight function μ on S , we define $(Bf)(\mu)$ as the weight function on S' defined by putting

$$(Bf)(\mu)(s') := \sum_{s \in f^{-1}(s')} \mu(s).$$

Similarly to the finitary distribution functor, the *finitary bag functor* B_ω is the restriction of B to bags with finite support, i.e., $B_\omega S := \{\mu : S \rightarrow \mathbb{N} \mid \sum_{s \in S} \mu(s) < \omega\}$.

(h) The *binary tree functor* is the functor $Id^2 := Id \times Id$.

There are various ways to obtain new functors from old.

Example B.7 Let F , F_0 and F_1 be set functors.

(a) The *composition* of F_0 and F_1 , denoted as $F_1 \circ F_0$, is defined in the obvious way, e.g. on objects we put $(F_1 \circ F_0)(S) := F_1(F_0(S))$.

(b) The *product* $F_0 \times F_1$ of F_0 and F_1 is given (on objects) by $(F_0 \times F_1)S := F_0S \times F_1S$, while for $f : S \rightarrow S'$, the map $(F_0 \times F_1)f$ is given as $((F_0 \times F_1)f)(\sigma_0, \sigma_1) := ((F_0f)(\sigma_0), (F_1f)(\sigma_1))$.

(c) The *co-product* $F_0 + F_1$ of F_0 and F_1 is defined in a similarly straightforward way (note that on **Set** we may think of co-product as disjoint union).

(d) Given a set D , we let the *D-exponent functor* F^D be defined as follows. Given a set S , we put $F^D(S) := (F(S))^D$, that is, the set of maps from D to FS . Given an arrow $f : S \rightarrow S'$ and a function $h : D \rightarrow FS$, we simply define the arrow F^Df as the function $(Ff) \circ h$.

Of specific interest in the context of coalgebra and modal logic is the following operation on set functors, which generalises the relation between Kripke frames and Kripke models to the level of arbitrary coalgebras over **Set**. Think of $\mathbf{Q}$ as an arbitrary but fixed set of *proposition letters*.

Definition B.8 Given a set functor T and a set $\mathbf{Q}$ of proposition letters, we define $T_{\mathbf{Q}}$ as the functor $T_{\mathbf{Q}} := K_{P\mathbf{Q}} \times T$. We may refer to $T_{\mathbf{Q}}$ as the *T-model functor* associated with $\mathbf{Q}$. $\triangleleft$

Definition B.9 The collection of *Kripke polynomial set functors* or *KPF* is defined by the following ‘grammar’:

$$K ::= K_C \mid Id \mid K_0 \times K_1 \mid K_0 + K_1 \mid K^D \mid P \circ K, \quad (74)$$

where C and D are sets. The *polynomial set functors* are the ones obtained by the same grammar without the powerset functor, and the *finitary KPFs* are obtained by the version of (74) where P is replaced with P_ω . $\triangleleft$

Apart from the operations described in Example B.7, in the context of coalgebras there is (at least) one other way of interest to obtain new set functors from old, namely, to take the *finitary version* of a functor. For this we need to introduce some notation and terminology concerning inclusions.

Definition B.10 Given two sets A, B with $A \subseteq B$, we let $\iota_B^A : A \rightarrow B$ denote the associated *inclusion map*, i.e., $\iota_B^A : a \mapsto a$; we will also write $f : A \hookrightarrow B$ to denote that $f = \iota_B^A$ (and so, in particular, this means that $A \subseteq B$). We say that a set functor T *preserves inclusions* if, for every pair of sets A, B with $A \subseteq B$, we have that $TA \subseteq TB$ and $T\iota_B^A = \iota_{TB}^{TA}$. $\triangleleft$

► **Examples:** power set vs mon nbh

In case $T : \mathbf{Set} \rightarrow \mathbf{Set}$ preserves inclusions we define the notion of *support* as follows.

Definition B.11 Let T be an inclusion preserving set functor, and suppose that $\alpha \in TA$. We say that a subset $B \subseteq A$ *supports* α if $\alpha \in TB$. $\triangleleft$

Definition B.12 Given a set functor T , we define the following operation T_ω on an arbitrary set S and an arbitrary function $f : S \rightarrow S'$:

$$\begin{aligned} T_\omega(S) &:= \{(T\iota_S^X)(\xi) \mid \xi \in TX \text{ for some } X \subseteq_\omega S\}, \\ T_\omega(f) &:= (Tf) \upharpoonright_{T_\omega S}. \end{aligned}$$

We will call T_ω the *finitary version* of T . $\triangleleft$

Given the definition of T_ω on functions, we may write Tf instead of $T_\omega f$ without causing confusion. Observe that we obtain $T_\omega S \subseteq TS$, and that, in case T preserves inclusions, the definition of $T_\omega S$ simplifies to $T_\omega S = \bigcup \{TX \mid X \subseteq_\omega S\}$. Whether this is the case or not, we always have that T_ω is a functor (which justifies that we call it the finitary *version* of T).

Proposition B.13 (1) If T is a set functor, then so is T_ω .
(2) If T preserves inclusions, then so does T_ω .

Proof. We only prove part (1) of the proposition, where the key property to establish is that for any map $f : S \rightarrow S'$, the arrow Tf is well-typed; that is, $\sigma \in T_\omega S$ implies $(Tf)\sigma \in T_\omega S'$. To see why this is the case, take an arbitrary object $\sigma \in T_\omega S$, and let $X \subseteq_\omega S$ and $\xi \in TX$ be such that $\sigma = (T\iota)(\xi)$, where we write $\iota = \iota_S^X$ to simplify notation.

Now consider the following two diagrams, where we let ι' denote the inclusion arrow $\iota' : f[X] \hookrightarrow S'$:

$$\begin{array}{ccc} X & \xrightarrow{\iota} & S \\ f|_X \downarrow & & \downarrow f \\ f[X] & \xrightarrow{\iota'} & S' \end{array} \qquad \begin{array}{ccc} TX & \xrightarrow{T\iota} & TS \\ T(f|_X) \downarrow & & \downarrow Tf \\ Tf[X] & \xrightarrow{T\iota'} & TS' \end{array}$$

Since the left diagram commutes, so does the right one. From this it follows that $(T_\omega f)(\sigma) = (Tf)(\sigma) = (Tf)(T\iota)(\xi) = (T\iota')(Tf|_X)(\xi)$ and since $f[X]$ is a *finite* subset of S' , with inclusion map ι' , this suffices to show that $(Tf)(\sigma) \in T_\omega S'$ indeed. QED

B.3 Limits and colimits in Set

Definition B.14 The *product* of two objects A_0 and A_1 in a category $\mathbf{C}$ is an object $A_0 \times A_1$, together with two *projection arrows* $\pi_i : A_0 \times A_1 \rightarrow A_i$, such that for any pair of arrows $f_i : X \rightarrow A_i$ there is a *unique* arrow $u : X \rightarrow A_0 \times A_1$ such that the following diagram commutes:

$$\begin{array}{ccccc} A_0 & \xleftarrow{\pi_0} & A_0 \times A_1 & \xrightarrow{\pi_1} & A_1 \\ & \searrow f_0 & \uparrow u & \nearrow f_1 & \\ & & X & & \end{array} \quad (75)$$

We will often denote the arrow u as $\langle f_0, f_1 \rangle$.

Dually, the *co-product* of two objects A_0 and A_1 in a category $\mathbf{C}$ is an object $A_0 + A_1$, together with two *insertion arrows* $\kappa_i : A_i \rightarrow A_0 + A_1$, such that for any pair of arrows $f_i : A_i \rightarrow X$ there is a *unique* arrow $v : A_0 + A_1 \rightarrow X$ such that the following diagram commutes:

$$\begin{array}{ccccc} A_0 & \xrightarrow{\kappa_0} & A_0 + A_1 & \xleftarrow{\kappa_1} & A_1 \\ & \searrow f_0 & \downarrow v & \swarrow f_1 & \\ & & X & & \end{array} \quad (76)$$

The mediating arrow v will usually be denoted as $[f_0, f_1]$. ◁

Example B.15 In the category **Set**, we take for the product of two sets S_0 and S_1 their *cartesian product* $S_0 \times S_1 := \{(s_0, s_1) \mid s_i \in S_i\}$, with the obvious projection maps $\pi_i : (s_0, s_1) \mapsto s_i$. For a concrete representation of the co-product of S_0 and S_1 we take the *sum* or *disjoint union* $S_0 + S_1 := (\{0\} \times S_0) \cup (\{1\} \times S_1)$, with the insertion maps $\kappa_i : s \mapsto (i, s)$.

Definition B.16 The binary product and co-product of Definition B.14 are easily generalised to (co-)products over an arbitrary index set I ; we omit the details, but introduce the notation $\prod_{i \in I} A_i$ and $\coprod_{i \in I} A_i$ for the product and co-product of the family $\{A_i \mid i \in I\}$. Products and co-products of the empty family are called *final* respectively *initial* objects of the category. ◁

Definition B.17 Given two ‘parallel’ arrows $f_i : A \rightarrow B$ in a category $\mathbf{C}$, we define an *equalizer* of f_0 and f_1 as an arrow $g : X \rightarrow A$ that satisfies the equality $f_0 \circ g = f_1 \circ g$, and the following condition. For every arrow $g' : X' \rightarrow A$ such that $f_0 \circ g' = f_1 \circ g'$, there is a *unique* arrow $u : X' \rightarrow X$ such that $g' = g \circ u$, cf. the diagram on the left:


$$\begin{array}{ccc} X & \xrightarrow{g} & A \xrightarrow[f_1]{f_0} B \\ \uparrow u & \nearrow g' & \\ X' & & \end{array} \quad \begin{array}{ccc} A \xrightarrow[f_1]{f_0} B & \xrightarrow{h} & Y \\ & \searrow h' & \downarrow v \\ & & Y' \end{array}$$

Dually, a *co-equalizer* of f_0 and f_1 is an arrow $h : B \rightarrow Y$ satisfying $h \circ f_0 = h \circ f_1$ and the universal property as indicated in the diagram to the right. ◁

Example B.18 In the category **Set**, as the equalizer of two functions $f_i : S \rightarrow S'$, we can take the set $\text{eq}(f_0, f_1) := \{s \in S \mid f_0 s = f_1 s\}$, together with the inclusion map $\iota : \text{eq}(f_0, f_1) \hookrightarrow S$.

As the co-equalizer of $f_0, f_1 : S \rightarrow S'$, consider the *quotient* S' under the equivalence relation E_{f_0, f_1} generated by the set $\{(f_0 s, f_1 s) \mid s \in S\}$, together with the quotient map $q : S' \rightarrow S'/E_{f_0, f_1}$ mapping any $s' \in S'$ to its own equivalence class under E_{f_0, f_1} .

Definition B.19 Given two arrows $f_i : A_i \rightarrow B$ in a category **C**, a *pullback* of f_0 and f_1 is an object P , together with two arrows $p_i : P \rightarrow A_i$ which satisfy $f_0 \circ p_0 = f_1 \circ p_1$, together with the following condition. Given any ‘competitor’ P' , with arrows $p'_i : P' \rightarrow A_i$ such that $f_0 \circ p'_0 = f_1 \circ p'_1$, there is a *unique* arrow $u : P' \rightarrow P$ such that $p'_i = p_i \circ u$, in a diagram:



Dually we define the notion of a *pushout* of two arrows $f_i : B \rightarrow A_i$. ◁

Example B.20 In the category **Set** we can define, given two functions $f_i : S_i \rightarrow S$, the set $\text{pb}(f_0, f_1) := \{(s_0, s_1) \in S_0 \times S_1 \mid f_0(s_0) = f_1(s_1)\}$, and show that this set, together with the projection maps $\pi_i : \text{pb}(f_0, f_1) \rightarrow S_i$, is the pullback of f_0 and f_1 .

For the pushout, consider arrows $f_i : A \rightarrow B_i$, for $i = 0, 1$. Let E be the equivalence relation on the disjoint union $B_0 + B_1$ which is generated from the set of pairs $\{(f_0 a_0, f_1 a_1) \mid a \in A\}$. Let C be the quotient of $B_0 + B_1$ under E , that is, C is the set of E -cells, and let $e_i : B_i \rightarrow C$, for $i = 0, 1$, be the quotient map, restricted to B_i . That is, $e_i(b_i)$ is the equivalence class of $b_i \in B_i$ under the equivalence relation E . It is not hard to see that C , together with the functions e_0 and e_1 , is the pushout of f_0 and f_1 . One may also show that e_i is injective (surjective) if f_i is so.

The concepts of product, equalizer and pullback are not independent.

Fact B.21 *The following are equivalent, for any category **C**:*

- (1) ***C** has finite products and equalizers;*
- (2) ***C** has pullbacks and a final object.*

More in general we can define the notion of a limit or colimit of a diagram.

Definition B.22 Let **J** and **C** be categories, and assume that **J** is *small*, that is, its collection of objects forms a set (rather than a proper class). A *diagram of type **J*** in **C** is a functor $D : \mathbf{J} \rightarrow \mathbf{C}$. We refer to **J** as the *index set* of D , and write D_i rather than $D(i)$, where i is an arbitrary object or *index* in **J**. ◁

Example B.23 Here are four examples of diagrams (where we do not draw identity arrows):

$$\begin{array}{c} A_0 \quad A_1 \end{array} \quad | \quad \begin{array}{c} A \xrightarrow{f_0} B \\ \xrightarrow{f_1} \end{array} \quad | \quad \begin{array}{c} A_0 \xrightarrow{f_0} B \\ \uparrow f_1 \\ A_1 \end{array} \quad | \quad \begin{array}{c} B_0 \\ \uparrow f_0 \\ A \xrightarrow{f_1} B_1 \end{array}$$

Definition B.24 A *cone* to a diagram $D : \mathbf{J} \rightarrow \mathbf{C}$ consists of an object C in $\mathbf{C}$, together with an arrow $c_j : C \rightarrow D_i$ for each object i in $\mathbf{J}$, such that for each arrow $e : i \rightarrow j$ in $\mathbf{J}$, the following diagram commutes:

$$\begin{array}{ccc} C & & \\ c_i \downarrow & \searrow c_j & \\ D_i & \xrightarrow{D(e)} & D_j \end{array} \tag{77}$$

A *morphism of cones* $\gamma : (C, c_i)_{i \in \mathbf{J}} \rightarrow (C', c'_i)_{i \in \mathbf{J}}$ is an arrow $\gamma : C \rightarrow C'$ in $\mathbf{C}$ such that $c_i = c'_i \circ \gamma$ for each index i :

$$\begin{array}{ccc} C & \xrightarrow{\gamma} & C' \\ c_i \searrow & & \downarrow c'_i \\ & & D_i \end{array}$$

The notion of a cone dualizes to that of a *co-cone* in the obvious way. We let $\mathbf{Cone}(D)$ and $\mathbf{CoCone}(D)$ denote the emerging categories of cones and co-cones, respectively. $\triangleleft$

Definition B.25 Let $D : \mathbf{J} \rightarrow \mathbf{C}$ be a diagram. A *limit* for D is a terminal object in the category $\mathbf{Cone}(D)$, and a *colimit* for D is an initial object in the category $\mathbf{CoCone}(D)$. A (co-)limit is called *finite* if the index category $\mathbf{J}$ is finite. $\triangleleft$

Spelled out, the limit of a diagram $D : \mathbf{J} \rightarrow \mathbf{C}$ is a cone to D , that is, an object C in $\mathbf{C}$, together with a family $p_i : C \rightarrow D_i$ of arrows in $\mathbf{C}$ satisfying the cone condition (77), and such that for any D -cone $(C', c'_i)_{i \in \mathbf{J}}$ there is a unique cone morphism $u : (C', c'_i)_{i \in \mathbf{J}} \rightarrow (C, c_i)_{i \in \mathbf{J}}$.

Example B.26 Limits for the first three diagrams in Example B.23 may easily be identified with respectively products, equalizers and pullbacks of the objects and morphisms displayed. Co-products, co-equalizers and pushouts can be seen as the co-limits of, respectively, the first, second, and fourth diagram.

Limits and colimits do not always exist; if limits (for a certain type of diagram) always exist in a category $\mathbf{C}$, we say that $\mathbf{C}$ *has limits (of that type)*.

Fact B.27 For any category $\mathbf{C}$ the following are equivalent.

- (1) $\mathbf{C}$ has all (finite) limits;
- (2) $\mathbf{C}$ has all equalizers and (finite) products.

Fact B.28 The category $\mathbf{Set}$ has all limits and colimits.

B.4 Properties of set functors

In this section we discuss some properties of set functors that are of interest in the setting of coalgebra. First we recall some definitions.

Definition B.29 Let T be a set functor.

- (1) T *restricts to finite sets* if TS is finite whenever S is finite.
- (2) T is *smooth* if it preserves weak pullbacks. $\triangleleft$

Proposition B.30 Let T be a set functor. Then T preserves surjections and non-empty injections. That is, if $f : A \rightarrow B$ is surjective then so is $Tf : TA \rightarrow TB$; and if $f : A \rightarrow B$ is injective then so is $Tf : TA \rightarrow TB$, provided A is non-empty.

Proof. First let $f : A \rightarrow B$ be surjective, and take any map $g : B \rightarrow A$ such that $f \circ g = \text{id}_B$. It follows by functoriality that $Tf \circ Tg = \text{id}_{TB}$ which immediately implies the surjectivity of Tf .

Similarly, if $A \neq \emptyset$ and $f : A \rightarrow B$ is injective then we may consider a map $h : B \rightarrow A$ such that $h \circ f = \text{id}_A$. Again by functoriality we find $Th \circ Tf = \text{id}_{TA}$ which implies the injectivity of Tf . QED

Remark B.31 The above proof breaks down for empty injections, and indeed one may come up with functors that do not preserve the injectivity of the empty map. For instance, as a variation of the constant functor K_C , consider a non-injective map $h : D \rightarrow C$. We let K_h be the set functor given as follows. On objects we have $K_h(\emptyset) := D$, while $K_h(A) := C$ if $A \neq \emptyset$; and for an arrow $f : A \rightarrow B$ we define $K_h(f)$ to be id_D if $A = B = \emptyset$, id_C if $A \neq \emptyset$, and h if $A = \emptyset \neq B$. In other words, for any non-empty set X we find $K_h(!_X) = h$. This problem is avoided by requiring the functor to be *standard*.

It is often convenient to assume that $TA \subseteq TB$ if $A \subseteq B$; the following property requires a bit more.

Definition B.32 Where A and B are sets such that $A \subseteq B$, we let $\iota_B^A : A \rightarrow B$ denote the corresponding *inclusion map*. We say that a set functor T is *standard* if it preserves inclusions, that is, if $A \subseteq B$ then $TA \subseteq TB$ and $T\iota_B^A = \iota_{TB}^{TA}$. $\triangleleft$

One benefit of this property is that working with standard functors may reduce cognitive (and notational) clutter. Note as well that, in the case where $A = \emptyset$, we find that $T\emptyset \subseteq TB$, and T maps the unique function $!_B : \emptyset \rightarrow B$ to the inclusion $\iota_{TB}^{T\emptyset}$. In other words, standard functors preserve the injectivity of all maps. The main advantage of standard functors is that they have various other nice properties. For instance, they preserve non-empty *intersections*.

Proposition B.33 Let T be a standard set functor. Then T preserves non-empty intersections, i.e., for any pair of sets A, B with $A \cap B \neq \emptyset$, we have

$$T(A \cap B) = TA \cap TB. \tag{78}$$

• If, in addition, T is smooth, then (78) holds for every pair of sets A, B .

Restricting attention to standard functors is relatively harmless in the setting of coalgebras, because of the following fact which states that set functors are ‘almost’ standard.

Proposition B.34 *Let T be some set functor. Then there is a standard set functor T' , and, for every set S , a bijection mapping T -coalgebra arrows $\sigma : S \rightarrow TS$ to T' -coalgebra arrows $\sigma' : S \rightarrow T'S$. This family of bijections induces an isomorphism between the categories $\text{Coalg}(T)$ and $\text{Coalg}(T')$.*

Proof. Following the proof of [A&T,III.4.5] we may find a standard functor T' such that the restrictions of T and T' to non-empty sets are naturally isomorphic. Let η denote this natural isomorphism; that is, we have a natural family of bijections $\eta_S : TS \rightarrow T'S$, for non-empty sets S . We extend η to $\text{Coalg}(T)$ as follows. Given a T -coalgebra $\mathbb{S} = (S, \sigma)$, we define the map $\sigma^\eta : S \rightarrow T'S$ by putting

$$\sigma^\eta := \begin{cases} !_{T'\emptyset} & \text{if } S = \emptyset \\ \eta_S \circ \sigma & \text{otherwise,} \end{cases}$$

and defining $\mathbb{S}^\eta := (S, \sigma^\eta)$. It is straightforward to verify that this constitutes the required isomorphism between $\text{Coalg}(T)$ and $\text{Coalg}(T')$. QED

Example B.35 The bag functor B of Example B.6 is not standard, but we may ‘standardize’ it by representing any map $\mu : X \rightarrow \mathbb{N}^\infty$ via its ‘positive graph’ $\{(x, \mu x) \mid \mu x > 0\}$. A similar approach works for the finitary bag functor B_ω , and for the distribution functors D and D_ω .

B.5 Relation lifting

Recall that the concept of *relation lifting* was introduced in Definition 3.6. For convenience we recall the definition.

Fix some set functor T . Given a relation $R \subseteq S_0 \times S_1$, consider R as a *span*

$$S_0 \xleftarrow{\pi_0} R \xrightarrow{\pi_1} S_1,$$

where $\pi_i : R \rightarrow S_i$ and $p_i : TS_0 \times TS_1 \rightarrow TS_i$ denote the respective projection maps. We define the *relation lifting* of R as the relation $\bar{T}R \subseteq TS_0 \times TS_1$ given by

$$\bar{T}R := \{((T\pi_0)(u), (T\pi_1)(u)) \mid u \in TR\}, \quad (79)$$

that is, $\bar{T}R$ is the image of TR under the map $\tau_R := \langle T\pi_0, T\pi_1 \rangle$. In a diagram:

$$\begin{array}{ccccc} TS_0 & \xleftarrow{T\pi_0} & TR & \xrightarrow{T\pi_1} & TS_1 \\ & \searrow p_0 & \downarrow \tau_R & \swarrow p_1 & \\ & & \bar{T}R & & \\ & & \downarrow & & \\ & & TS_0 \times TS_1 & & \end{array}$$

In the propositions below we list various properties of relation lifting — proof details can be found in [KKV]. We start with some properties of $\bar{T}$ that hold for any functor T .

Fact B.36 Let T be a set functor. Then the relation lifting $\bar{T}$ has the following properties:

1. $\bar{T}$ extends T : $\bar{T}(\text{Gr}f) = \text{Gr}(Tf)$ for all functions $f : A_0 \rightarrow A_1$,
2. $\bar{T}$ preserves diagonals: $\bar{T}\Delta_A = \Delta_{TA}$ for any set A ;
3. $\bar{T}$ preserves converse: $\bar{T}R^\sim = (\bar{T}R)^\sim$ for all relations $R \subseteq A_0 \times A_1$;
4. $\bar{T}$ is monotone: $R \subseteq Q$ implies $\bar{T}R \subseteq \bar{T}Q$ for all relations $R, Q \subseteq A_0 \times A_1$;
5. $\bar{T}$ is semi-functorial: $\bar{T}(R; Q) \subseteq \bar{T}R; \bar{T}Q$, for all relations $R \subseteq A_0 \times A_1, Q \subseteq A_1 \times A_2$.

In case the functor T is *smooth*, Proposition B.36(5) can be strengthened to full functoriality. In fact, we have the following equivalence, which explains the importance of smoothness in the theory of coalgebras.

Fact B.37 For any set functor T the following are equivalent:

- (1) T is smooth;
- (2) $\bar{T}(R; Q) = \bar{T}R; \bar{T}Q$, for all relations $R \subseteq A_0 \times A_1, Q \subseteq A_1 \times A_2$.
- (3) $\bar{T}$ is an endofunctor on the category Rel of sets and binary relations.

•

In case T is standard, some more properties are preserved.

Fact B.38 Let T be a smooth and standard set functor. Then the relation lifting $\bar{T}$ has the following properties:

1. $\bar{T}$ preserves domain: $\text{Dom}(\bar{T}R) = T(\text{Dom}R)$, for all relations $R \subseteq A_0 \times A_1$;
2. $\bar{T}$ preserves range: $\text{Ran}(\bar{T}R) = T(\text{Ran}R)$, for all relations $R \subseteq A_0 \times A_1$;
3. $\bar{T}$ preserves composition: $\bar{T}(R; Q) = \bar{T}R; \bar{T}Q$, for all relations $R \subseteq A_0 \times A_1, Q \subseteq A_1 \times A_2$;
4. $\bar{T}$ preserves restrictions:

$$\bar{T}(R \upharpoonright_{B_0 \times B_1}) = (\bar{T}R) \upharpoonright_{TB_0 \times TB_1}$$

for all relations $R \subseteq A_0 \times A_1$, and all sets $B_0 \subseteq A_0$ and $B_1 \subseteq A_1$;

5. $\bar{T}_\omega$ aligns with $\bar{T}$: $\bar{T}_\omega R = \bar{T}R \cap (T_\omega A_0 \times T_\omega A_1)$, for all relations $R \subseteq A_0 \times A_1$.